Hull–White Short-Rate Models

Martin Burke (with Claude)

Draft

Abstract

Modelling interest rates is harder than modelling a stock, and the difference is structural, not merely technical. A stock is one number; a rate market is an entire *curve*—a discount factor for every maturity—which the model must move coherently, which performs its own discounting, and which it must match *exactly* today or admit arbitrage. This note builds the Hull–White model from that starting point. It opens by starting from the familiar equity-derivatives picture—a single underlying—and stepping to a whole term structure, and the requirements that leap imposes—then meets them the way a short-rate model does: collapse the curve onto one mean-reverting rate, and fit today's prices with a free, time-dependent drift. Part I develops the one-factor model—the Gaussian short rate, the drift that fits the curve, the affine bond-price formula, the vanilla options it prices in closed form (with a worked caplet), how its two parameters are calibrated, the smile it *cannot* make, and the early-exercise exotics it values on a lattice—before exposing the one thing a single factor cannot do: move rates at different maturities by anything other than lockstep. Part II adds a second factor (the G2++ model) so the curve can twist as well as shift, decorrelating rates and fitting a humped volatility term structure a single factor misses. Throughout, the spine is a two-layer split—a drift that nails today's curve exactly, and a few volatility parameters left for the dynamics—and a diagram for every idea.

Contents

1 From one number to a whole curve 3

2 The objects, quickly 5

3 Short-rate models and the central tension 7

I The one-factor model 9

4 The one-factor Hull–White model 9

5 Fitting today's curve 12

6 From the short rate to bond prices 15

7 What the one-factor model prices: the vanillas 18

8 Calibrating the one-factor model 24

9 The smile the model makes, and the skew it cannot 28

10 Beyond closed form: Bermudans and callables 30

11 What one factor cannot do 33

II The two-factor model 33

12 Why a second factor 33

13 The two-factor (G2++) model 35

14 Pricing and calibration with two factors 37

15 What two factors price that one cannot 39

16 Choosing a model, and what lies beyond 41

A From the short rate to the bond price 42

B The Hull–White trinomial tree 42

C The forward measure and the Black formula 43

D Anatomy of a vanilla swap and swaption 46

E Jamshidian's decomposition for swaptions 48

F Solving the OU process: mean and variance of r_t 49

1 From one number to a whole curve

Start from the familiar picture in equity derivatives. There a model has one underlying: a spot price S_t, a single number that wanders up and down (Figure 1, left). A whole volatility surface may sit on top of it, but the thing being modelled—the state that an equation like $\mathrm{d}S_t = \mu S_t\,\mathrm{d}t + \sigma S_t\,\mathrm{d}W_t$ pushes forward—is that one number. Interest rates break this comfortable picture, and it is worth being precise about how, because every design choice in the rest of this note is a response to the difference.

In a rates market there is no single underlying. The object you must model is an entire *curve*: the price today of a zero-coupon bond $P(t,T)$ for every maturity T—equivalently, an interest rate for every horizon (Figure 1, right). At each instant the state is not a point but a whole function of maturity, and as time passes that function moves: it is a curve that shifts, steepens and wriggles. Modelling a wandering number is one thing; modelling a wandering *curve* is the leap. Four complications follow from it, and naming them now tells us exactly what a rates model must do.

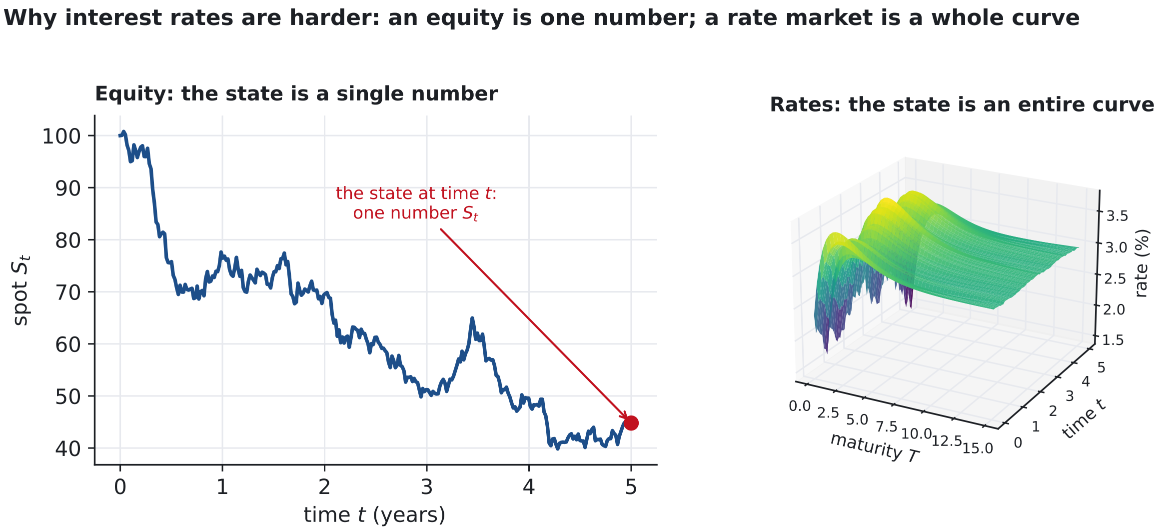

Figure 1: The leap. *Left:* an equity model evolves a single number, the spot S_t. *Right:* a rates model evolves an entire curve—a rate for every maturity—and that whole curve moves as time passes (shown as a surface over maturity and time). At any instant the curve is smooth across maturity, but its evolution *along* time is jagged, because rates move randomly. The state is no longer a point but a function; everything hard about interest-rate modelling starts here.

1. The curve's points are not free—they are arbitrage-linked. The maturities cannot move independently. Today's one-year and two-year bonds already imply the one-year interest rate, one year forward; if that implied forward drifted out of line with the traded one-year bond, you could borrow at one rate and lend at the other for a riskless profit. So the points of the curve are bound together by no-arbitrage (Figure 2): you cannot move one maturity without the others cohering. A rates model must keep *every* such relation, across all maturities, at once—a constraint a lone equity spot, being one number, never faces.

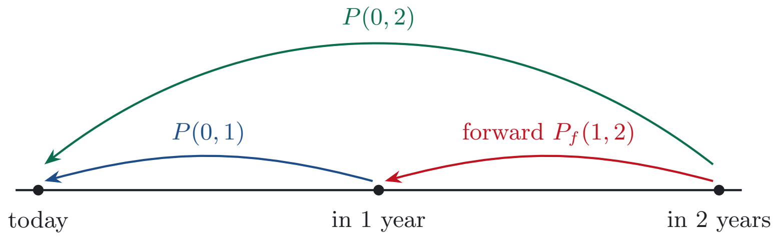

$P(0,2) = P(0,1) \times P_f(1,2)$ — fix any two, and the third is forced.

Figure 2: No-arbitrage binds the maturities together. Today's one- and two-year bonds already imply the one-year rate one year forward; were the implied forward to disagree with the traded one-year bond, the gap would be a riskless profit. A rates model must honour every such link, across all maturities, simultaneously—something a single equity spot never has to.

2. The underlying is also the discounting. In equity pricing you discount a payoff to today at an interest rate taken as given—a separate, usually deterministic input. In a rates market the curve you are modelling *is* that discounting. The random object and the discount factor are one and the same (Figure 3): every future cashflow—including the bonds the model must reprice—is pulled back to today along the very curve whose randomness you are trying to capture. The model feeds on its own output, a circularity equity pricing never confronts.

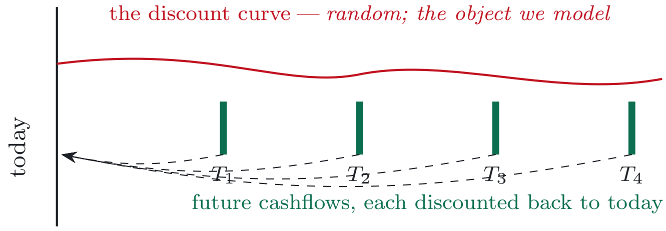

Figure 3: The underlying is the discounting. Every future cashflow is pulled back to today (dashed arrows) along the discount curve (red). In equities that curve is a fixed input; in rates it is exactly the random object being modelled—so the model must discount its own cashflows along its own output.

3. Fitting today's prices is non-negotiable. An equity model that misprices a vanilla by a little is wrong but usable, and the underlying spot is simply observed—one number, nothing to fit. A rates model is held to a sharper standard. The bonds are the most liquid instruments in the market, and they *are* the curve; a model that does not reproduce $P(0,T)$ exactly would value a traded bond away from its price—an arbitrage against the very instruments it hedges with, and an inconsistent discounter for everything else. Matching today's whole curve exactly is therefore the first, non-negotiable requirement, in a way that has no real equity analogue.

4. The dynamics are different, and richer. A stock is expected to grow; a rate is tethered. Central banks, inflation and the business cycle pin interest rates to a range, so they *mean-revert* rather than wander off—a qualitatively different process from a stock's exponential drift. And because the state is a whole curve, its points move with a rich *correlation structure*: near and far rates move together, but not identically, and capturing how they decouple is part of the job. Where the equity quant had a lone scalar, the rates quant has a structured, mean-reverting, internally correlated curve.

Taming the curve with one number. Four complications, one curve. The escape that makes the whole subject tractable is disarmingly simple: rather than model the infinite-dimensional curve directly, model a *single* quantity—the *short rate* r_t, the interest rate for the very next instant—and *derive* the entire curve from it by no-arbitrage (Figure 4). One number, evolved through one stochastic differential equation, regenerates a bond price for every maturity. This is the *short-rate* approach, and it collapses the curve back to something the equity quant would recognise: a single state variable. The price of that collapse—that all maturities are now driven by one number—is a story we return to (Section 11); for now it is the key that unlocks everything.

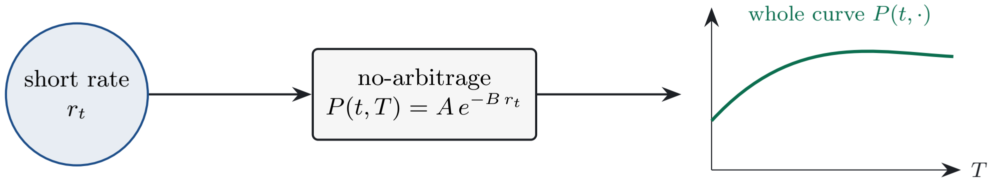

Figure 4: The short-rate idea. One number r_t, passed through the no-arbitrage pricing relation, regenerates a bond price for every maturity—the whole curve. An infinite-dimensional object is tamed back to a single state variable, at the cost (Section 11) that one number now drives every point of the curve.

The four complications then crystallise into two concrete demands. The model must (i) *reprice today's whole curve exactly* (complications 1–3), and (ii) *move that curve realistically* (complication 4). Hold those two requirements in mind: the entire construction of Hull–White is the cleanest known way to meet both at once. The next two sections pin down the objects (§2), then turn the two requirements into the model (§3).

2 The objects, quickly

This section fixes notation for the handful of objects the rest of the note uses; anyone who has priced in rates will have met them, so it is a recap, not a primer.

Everything derives from the *discount factor*, equivalently the zero-coupon bond price, $P(t,T)$: the value at time t of one unit of currency paid for certain at T. A few rates re-express the same information (Figure 5). The *zero rate* $R(t,T)$ is the single constant rate that discounts T back to t, via $P(t,T) = e^{-R(t,T)(T-t)}$—an average rate over the whole horizon.

The *forward rate* is the rate you can lock in *today* for borrowing or lending over a *future* window $[T_1, T_2]$. It is not a free quote: no-arbitrage pins it to the two bonds straddling the window. Agreeing now to lend \$1 from T_1 to T_2 must cost the same whether you arrange it directly or build it synthetically—buy one T_2-bond and fund it by shorting $P(t,T_2)/P(t,T_1)$ of the T_1-bond—so the simply-compounded forward rate is forced to be

$$F(t;T_1,T_2) = \frac{1}{T_2 - T_1}\left(\frac{P(t,T_1)}{P(t,T_2)} - 1\right). \tag{1}$$

This is exactly the linkage of Figure 2: the forward is whatever keeps the near bond, the far bond and the forward loan mutually consistent. Shrinking the window to an instant ($T_2 \to T_1 = T$) gives the *instantaneous forward rate*

$$f(t,T) = -\frac{\partial}{\partial T}\ln P(t,T), \tag{2}$$

the rate for an infinitesimal loan starting at T—a marginal rate, the leading edge of the curve. The zero rate is the running average of these forwards, $R(t,T) = \frac{1}{T-t}\int_t^T f(t,u)\,\mathrm{d}u$, which is why the forward curve is the steeper of the two.

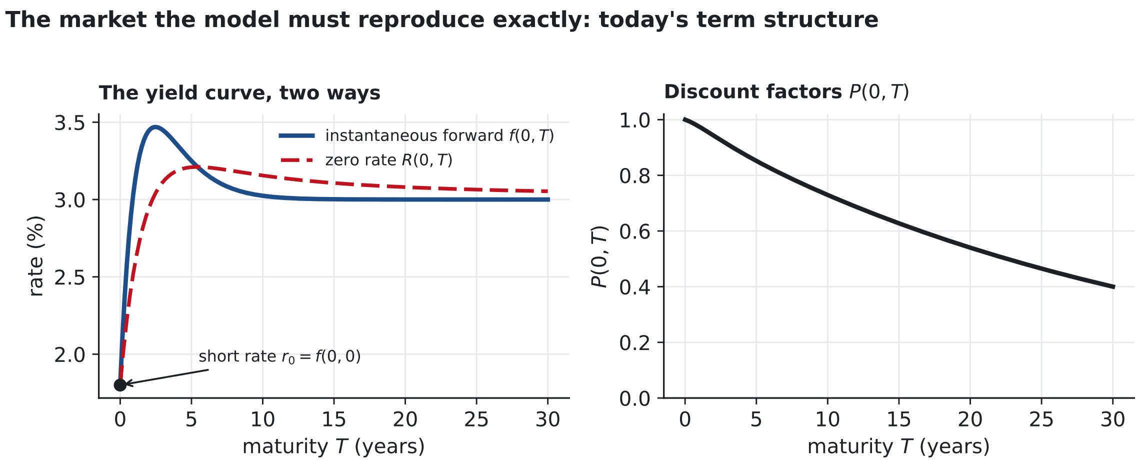

Figure 5: Today's term structure, three equivalent views. *Left:* the same curve as the instantaneous forward rate $f(0,T)$ (the marginal rate for a loan starting at T) and as the zero rate $R(0,T)$ (its running average to T); the forward is the steeper, leading object. The *short rate* is the very front end, $r_0 = f(0,0)$ (marked)—the rate for the next instant, and the one quantity the model evolves. *Right:* the discount factors $P(0,T) = e^{-\int_0^T f(0,s)\,\mathrm{d}s}$ those rates imply. Everything in this note is built on one smooth synthetic curve of this shape.

The *short rate* is the front end of this curve—the instantaneous forward for the very next instant, $r_t = f(t,t)$, the rate on an overnight loan. It is the one quantity a short-rate model evolves, and from it no-arbitrage rebuilds the entire curve.

Why one number prices the whole curve: the deterministic case. Take the easy world first, where the short rate follows a *known*, deterministic path r_s. Then there is no uncertainty in the discounting, only the time value of money. Roll \$1 at the short rate, reinvesting instant by instant, and by time t it has grown into the *money-market account*

$$\beta_t = \exp\Big(\int_0^t r_s\,\mathrm{d}s\Big), \qquad \text{so over } [t,T] \text{ it grows by } \frac{\beta_T}{\beta_t} = \exp\Big(\int_t^T r_s\,\mathrm{d}s\Big). \tag{3}$$

To have \$1 for certain at T you therefore set aside the reciprocal of that growth factor today. A T-bond pays exactly that \$1, so its price *is* that amount:

$$\begin{aligned} P(t,T) &= \frac{\beta_t}{\beta_T} \\ &= \exp\Big(-\int_t^T r_s\,\mathrm{d}s\Big) \\ &= \exp\Big(-\int_t^T f(t,u)\,\mathrm{d}u\Big) \qquad \text{(deterministic rates).} \end{aligned} \tag{4}$$

The last equality is no coincidence: with rates known, the instantaneous forward (2) is just the future short rate itself, $f(t,u) = r_u$, so discounting at the short rate and discounting along the forward curve are the same thing. One known path of r already gives the whole curve.

Restoring the randomness. Now let r_s be stochastic. The rolled-up amount $\int_t^T r_s\,\mathrm{d}s$ is no longer known at time t, so the discount factor $\exp(-\int_t^T r_s\,\mathrm{d}s)$ is itself a *random variable*, and (4) cannot hold path by path—the bond has one definite price today, but the realised discount factor will only be known at T. The bond's price must therefore be some *average* of that random discount factor; the only question is which. Not the naive real-world expectation, which would ignore the premium investors demand for bearing rate risk. No-arbitrage answers cleanly through the same money-market account (3)—the one riskless self-financing strategy available to everyone—under which *every* asset, measured in units of β, has no predictable drift: it is a martingale under a particular *risk-neutral* measure $\mathbb{Q}$. Imposing that on the T-bond, $P(t,T)/\beta_t$ is a $\mathbb{Q}$-martingale, and rearranging gives the deterministic formula with an expectation wrapped around it,

$$\begin{aligned} P(t,T) &= \beta_t\, \mathbb{E}^{\mathbb{Q}}\Big[\beta_T^{-1} \,\big|\, \mathcal{F}_t\Big] \\ &= \mathbb{E}^{\mathbb{Q}}\left[\exp\Big(-\int_t^T r_s\, \mathrm{d}s\Big) \,\Big|\, \mathcal{F}_t\right]. \end{aligned} \tag{5}$$

The intuition is unchanged—a no-arbitrage tie between the two ways of holding cash to T, buy the bond or roll overnight—now read in expectation: the bond is worth what the rolling strategy is *expected* to deliver, in the risk-adjusted world where every asset earns the short rate on average. And when the rate happens to be known, the expectation falls away and (5) collapses back to the deterministic (4). Either way, fixing the dynamics of the single number r_t fixes the whole curve, every $P(t,T)$ at once: the bridge promised in §1. The art is in making that curve come out right.

3 Short-rate models and the central tension

The relation (5) is what makes the short-rate idea work, but it also sets up the central difficulty, because it asks two things of the same small object r_t at once. Through the *average* of its future path, r_t must reproduce today's curve exactly—requirement (i). Through the *fluctuations* of that path, it must generate realistic movements of the curve—requirement (ii). These pull in different directions, and the history of short-rate modelling is the story of meeting both (Figure 6).

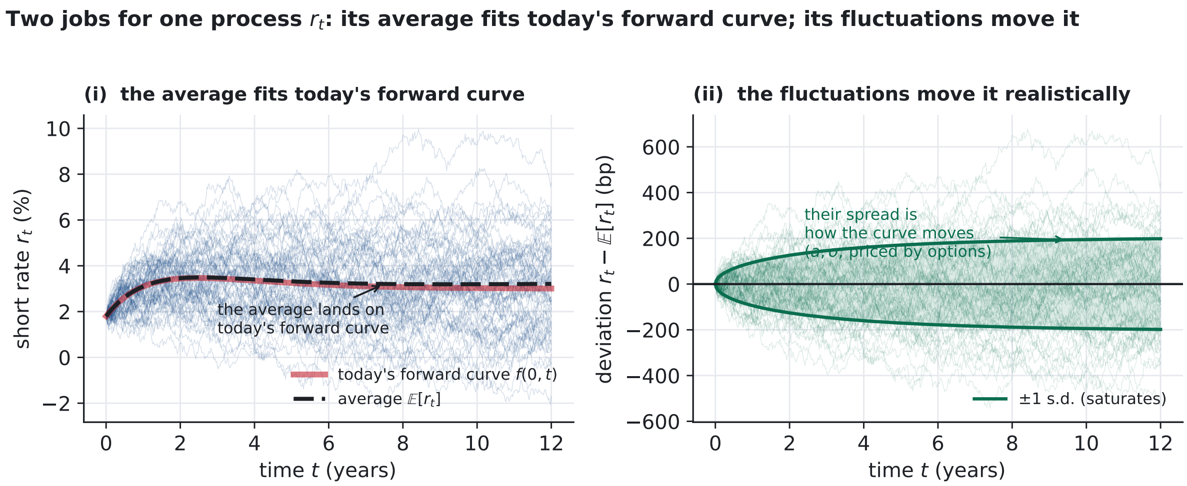

Figure 6: The same process r_t is asked to do two jobs at once. **(i)** The *average* of its future path reproduces today's *forward* curve $f(0,t)$—the dashed mean lands on the red curve (exactly, once the small convexity correction of §5 is included). **(ii)** The *fluctuations* of that path, whose spread is set by the volatility parameters and saturates under mean reversion, are what move the curve over time—and what the prices of interest-rate options see. Requirement (i) is a constraint on the mean; requirement (ii) is a constraint on the spread; the art is to satisfy both without one stealing from the other.

What does requirement (ii) actually mean? As calendar time passes, today's curve does not stay put: tomorrow's is a shifted, re-shaped version of it, and over weeks and months it steepens, flattens, or moves up and down bodily. In the model every one of those movements comes from the random shocks $\sigma\, \mathrm{d}W_t$ buffeting r_t. *Realistic* means matching those movements to a benchmark—the movements the market itself expects, made concrete by the prices of interest-rate options (caps, floors, swaptions), and corroborated by how the curve has actually behaved in the past. There are two things to match. The *size*: how many basis points a given rate typically moves in a year, which is exactly what an option's implied volatility measures. And the *shape*: short-maturity rates moving more than long-maturity ones, and (ideally) different maturities moving only partly in step rather than rigidly together—a last refinement that, as we will see, a single short rate cannot manage (§11) and that motivates the second factor of Part II. Both the size and the shape are read off the same option prices, because an interest-rate option is a bet on how far a rate will move; so requirement (ii) is, in practice, the demand that the model reproduce the volatilities those prices imply—the magnitudes maturity by maturity, and the way different maturities co-move. Keep this split in mind—the average against the curve, the fluctuations against option prices—because it is the division of labour the rest of the note is built on.

Vasicek: the simplest mean-reverting short rate. The first short-rate models could not. Vasicek [1] proposed the simplest mean-reverting diffusion for the short rate,

$$\mathrm{d}r_t = a\,(b - r_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{6}$$

governed by three *constants*: a long-run level b, a reversion speed a, and a volatility σ. Read the drift as a spring (a picture we return to in §4): wherever the rate is, it is pulled back toward b at a rate set by a, while σ shakes it. Because the coefficients are constant, the equation can be solved in closed form with the integrating factor e^{at}:

$$r_t = b + (r_0 - b)\,e^{-at} + \sigma \int_0^t e^{-a(t-s)}\,\mathrm{d}W_s, \tag{7}$$

from which the first two moments follow at once,

$$\begin{aligned} \mathbb{E}[r_t] &= b + (r_0 - b)\,e^{-at}, \\ \operatorname{Var}[r_t] &= \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big). \end{aligned} \tag{8}$$

Both make the spring quantitative. The mean slides from today's rate r_0 to the long-run level b, closing the gap at the exponential rate a. The variance, rather than growing without bound as a random walk's would, *saturates* at $\sigma^2/2a$: mean reversion pulls the rate back the harder the further it strays, holding it inside a band whose width is set by the contest between the shock size σ and the pull a. (Appendix F carries out this integrating-factor solution and these two moments in full, for the mean-zero version of the process that underlies the Hull–White formulae.) Cox–Ingersoll–Ross [2] kept the same constant-coefficient drift but scaled the volatility by $\sqrt{r_t}$, so the rate cannot go negative. Both are economically sensible, but both share one fatal limitation: with only a handful of constants, the curve they *produce* through (5) is a single fixed shape, and it will not line up with the curve the market actually shows (Figure 5). They meet requirement (ii) tolerably and fail requirement (i) outright—and failing (i), as we saw, is fatal: a model that misprices today's bonds is arbitrageable against the most liquid instruments there are.

Hull–White's fix: free the drift in time. Hull and White [3] fixed this with one surgical change to Vasicek's equation (6). Multiply out its drift, $a(b - r_t) = ab - a\,r_t$: the only thing tying the model to a fixed curve is the *constant* ab. Hull–White promotes that constant to a free function of time,

$$\underbrace{ab}_{\text{Vasicek: a constant}} \quad \longrightarrow \quad \underbrace{\theta(t)}_{\text{Hull–White: a free function}},$$

so the rate now reverts toward a *time-dependent* target $\theta(t)/a$—the defining feature of the one-factor Hull–White short rate, whose equation (9) we write out and unpack in §4. (We use $\theta(t)$ freely in the rest of this section before that equation appears; it is exactly Vasicek's ab set loose in time.) The crucial point is that $\theta(t)$ is not a parameter to be guessed but a function *solved for*, so that (5) reproduces the entire initial curve by construction. The volatility structure stays exactly as simple as Vasicek's—one speed a, one volatility σ—and these stay genuinely free, to be set by how the curve should *move*.

The two-layer split. This is the model's defining division of labour, and it is the spine of everything that follows (Figure 7):

- an *exact layer*—the time-dependent drift—that pins today's curve by construction, for *any* values of the volatility parameters;
- a *dynamics layer*—the mean reversion a and volatility σ—left free to match how the curve moves, which is what the prices of interest-rate options see.

The two layers answer to two different markets and never compete: the drift takes care of requirement (i), exactly, leaving the handful of dynamics parameters entirely to requirement (ii). (Readers who know local-stochastic-volatility models will recognise the pattern—a deterministic device that absorbs the static fit exactly, freeing a few parameters for the dynamics; there the device is a leverage function fitting the vanilla surface, here a drift fitting the discount curve.)

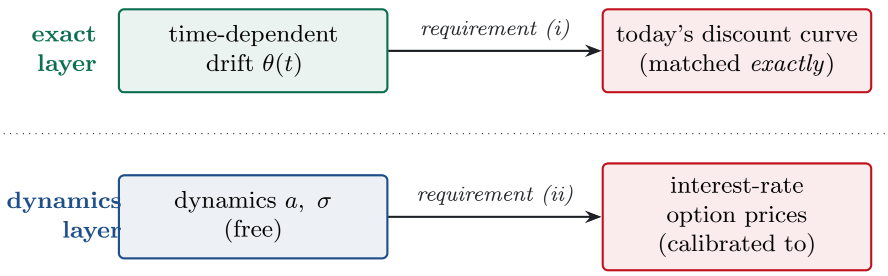

Figure 7: The spine of the note. The time-dependent drift (top) handles requirement (i)—fitting today's curve exactly—for any volatility parameters; the two dynamics parameters a, σ (bottom) handle requirement (ii)—realistic movement—calibrated to option prices. The layers answer to different markets and never compete.

Part I now makes each layer concrete for the one-factor model: the short-rate equation and its dynamics (§4), the drift that fits the curve (§5), the bond-price formula that turns the one number back into the whole curve (§6), the vanilla options it prices in closed form (§7), how its two parameters are calibrated to those prices (§8), the smile it makes and the skew it cannot (§9), and the path-dependent and early-exercise products that need a lattice (§10)—before we reach the one thing a single factor cannot do.

Part I
The one-factor model

4 The one-factor Hull–White model

The one-factor Hull–White short rate solves the stochastic differential equation

$$\mathrm{d}r_t = \big(\theta(t) - a\,r_t\big)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{9}$$

under the risk-neutral measure. This is exactly Vasicek's equation (6) with the one change made in §3: the constant drift term ab has become the free function $\theta(t)$. Read the drift as a spring: wherever r_t is, it is pulled back toward the level $\theta(t)/a$ at a speed set by the *mean-reversion rate* $a > 0$. The *volatility* σ is the size of the random shocks. Two features of this equation are worth dwelling on, because they are the whole character of the model.

It is mean-reverting. Unlike a stock, which diffuses ever further from its start, the short rate is tethered. Push it up and the negative $-a\,r_t$ term pulls it back; the higher it goes the harder the pull. The consequence is visible in the spread of the rate over time (Figure 8): the cloud of possible future rates fans out at first, just like a diffusion, but then *stops* widening. Its variance saturates at $\sigma^2/2a$ rather than growing without bound. This is exactly the behaviour real rates show—they do not wander off to arbitrary levels—and the speed a controls how quickly the fan reaches its stable width.

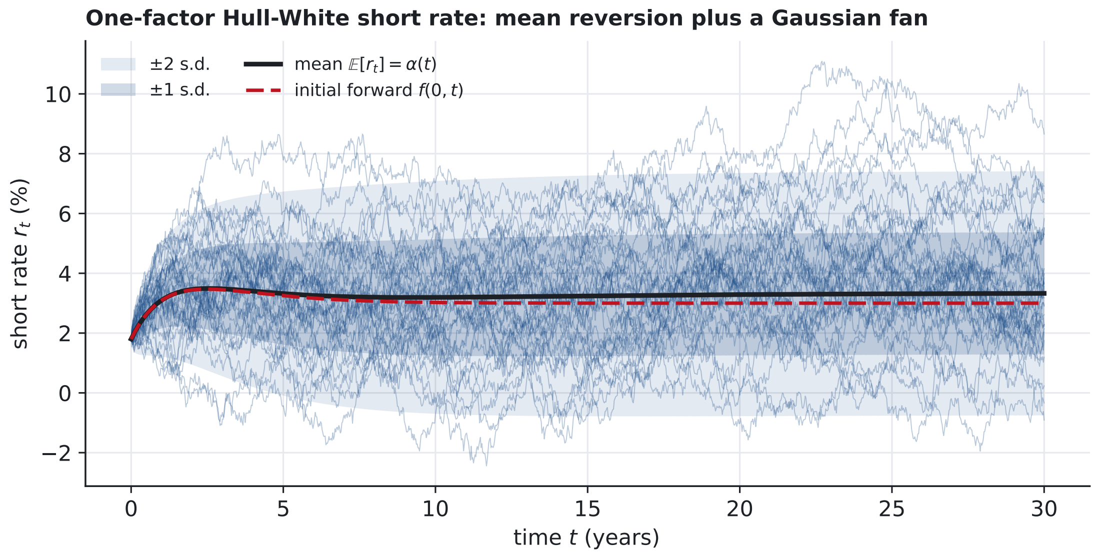

Figure 8: Simulated one-factor Hull–White short-rate paths (thin blue), with the mean $\mathbb{E}[r_t] = \alpha(t)$ (black) and $\pm 1, 2$ standard-deviation bands. The fan widens then stabilises—mean reversion caps the spread—and the mean tracks the initial forward curve $f(0, t)$ (red dashed) up to a small convexity gap that the model adds deliberately (Section 5). Here $a = 0.12$, $\sigma = 1\%$.

It is Gaussian. The shock term $\sigma\,\mathrm{d}W_t$ has a constant coefficient: it does not depend on the level of r_t. That makes (9) a linear SDE, and its solution is *normally distributed*. Every r_t is Gaussian, every bond price is a lognormal-type expression, and a great many prices come out in closed form. The price of this convenience is that a Gaussian rate can go *negative*: there is nothing in the model to stop it. For decades this was filed under "known defect"; since several major currencies spent years at negative policy rates, it looks more like a feature. Where strictly positive rates are wanted, one moves to a model with a $\sqrt{r}$ or shifted diffusion (CIR, or the shifted "CIR++"), at the cost of some tractability.

A cleaner way to hold the model: $r = x + \alpha$. The time-dependent drift $\theta(t)$ makes (9) look more complicated than it is. The tidy reformulation, due to the structure of the equation, is to split the short rate into a deterministic part and a random part,

$$r_t = x_t + \alpha(t), \qquad \mathrm{d}x_t = -a\,x_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \quad x_0 = 0, \tag{10}$$

where x_t is a plain mean-reverting Gaussian process reverting to *zero*—an Ornstein–Uhlenbeck process carrying all the randomness and none of the curve-fitting—and $\alpha(t)$ is a deterministic shift carrying all the curve-fitting and none of the randomness (Figure 9). All the dynamics live in x_t; all the calibration to today's curve lives in $\alpha(t)$. This is the same exact-layer / dynamics-layer split as before, now visible in the state variable itself, and it is how the model is almost always implemented: simulate the simple zero-mean x_t, then add the precomputed shift.

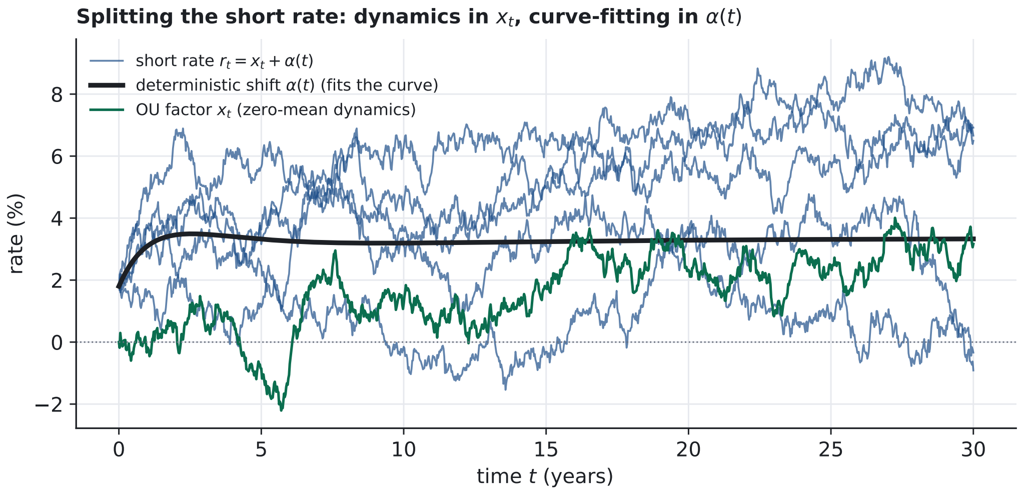

Figure 9: The decomposition $r_t = x_t + \alpha(t)$. The deterministic shift $\alpha(t)$ (black) is a fixed curve that pins the model to today's term structure; the Ornstein–Uhlenbeck factor x_t (green) is a zero-mean mean-reverting process carrying the randomness. Their sum is the short rate (blue). Simulate the simple part, add the fixed part.

Because x_t has mean zero, the mean of the short rate is just the shift, $\mathbb{E}[r_t] = \alpha(t)$, and because x_t is Ornstein–Uhlenbeck its variance is the saturating expression (derived in full in Appendix F)

$$\begin{aligned} \mathbb{E}[r_t] &= \alpha(t), \\ \mathrm{Var}[r_t] &= \frac{\sigma^2}{2a}\left(1 - e^{-2at}\right) \xrightarrow[t\to\infty]{} \frac{\sigma^2}{2a}, \end{aligned} \tag{11}$$

the limit being the stationary variance the fan settles into. The speed a thus does double duty: it sets how fast a shock decays (a disturbance to x_t fades like e^{-at}) and, through (11), how wide the long-run distribution of rates is. We return to this when calibrating (Section 8). Figure 10 plots the two moments against the horizon.

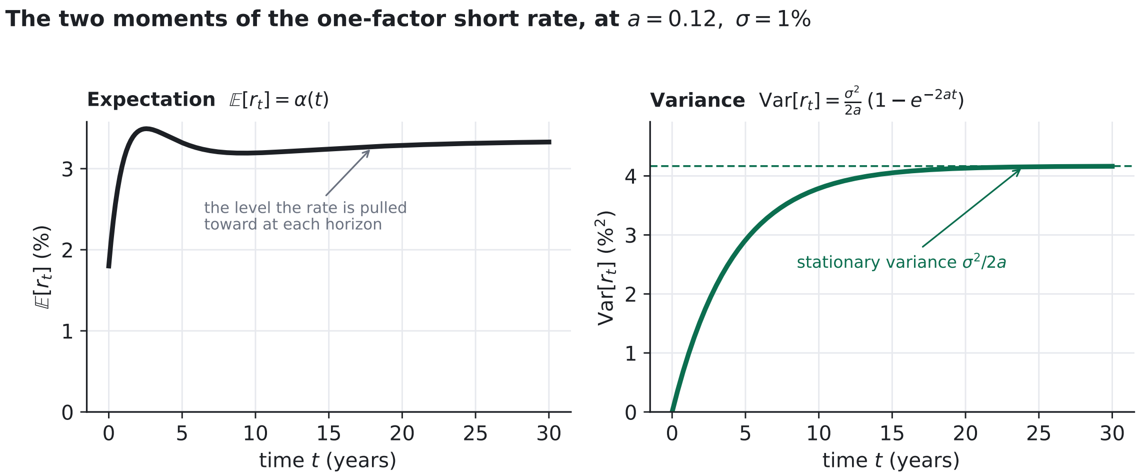

Figure 10: The two moments of (11), as functions of the horizon t, for the note's base case $a = 0.12$, $\sigma = 1\%$. *Left:* the expectation $\mathbb{E}[r_t] = \alpha(t)$—the deterministic level the rate is pulled toward at each horizon (it traces today's forward curve, up to the small convexity term of §5). *Right:* the variance $\mathrm{Var}[r_t]$, which starts at zero, grows, and *saturates* at the stationary value $\sigma^2/2a$ (dashed)—the signature of mean reversion, where a random walk's variance would instead grow without bound. The two together are the entire one-dimensional distribution of r_t, since it is Gaussian.

5 Fitting today's curve

The time-dependent drift exists for one reason: to make the model reprice the initial curve. There are two equivalent ways to pin it down: through the shift $\alpha(t)$ in the clean decomposition $r = x + \alpha$, or through the time-dependent drift $\theta(t)$ in the original SDE (9). The shift is by far the more transparent, so we lead with it and turn to the drift form afterwards.

The clear way: fitting the shift $\alpha(t)$

The matching logic. Work in the decomposition $r_t = x_t + \alpha(t)$ of (10). The Ornstein–Uhlenbeck factor x_t is completely fixed—zero mean, reverting to zero, carrying all the randomness—so the *only* freedom left for fitting today's curve is the deterministic shift $\alpha(t)$. The model's time-zero bond prices are $P^{\mathrm{mod}}(0,T) = \mathbb{E}\big[\exp(-\int_0^T r_s\,\mathrm{d}s)\big]$, and because the model is Gaussian and affine these have a closed form (Appendix A). Fitting the curve means choosing α so that

$$P^{\mathrm{mod}}(0,T) = P(0,T) \quad \text{for } \textit{every} \text{ maturity } T.$$

That looks like infinitely many equations—one per T—for the single unknown function $\alpha(\cdot)$, but they line up exactly. Take logarithms and differentiate in T: because the instantaneous forward rate is $f(0,T) = -\partial_T \ln P(0,T)$, the condition turns into "the model's forward curve must equal the market's", one equation at each maturity, and the forward at T depends on α only up to time T. So the equations solve *pointwise*, sweeping T outward from zero, with no optimisation anywhere. Carrying this out (Appendix A) gives a closed form for the shift,

$$\alpha(t) = \underbrace{f(0,t)}_{\text{initial forward}} + \underbrace{\frac{\sigma^2}{2a^2}\big(1 - e^{-at}\big)^2}_{\text{convexity term}}. \tag{12}$$

Reading the two terms. The shift splits into a leading piece and a small correction (Figure 11). The *leading* piece is just today's forward curve, $\alpha(t) \approx f(0,t)$: to a first approximation the model sets the average future short rate equal to the forward rate the market already quotes for that date—which is what no-arbitrage demands, since the forward rate is essentially the market's own expectation of the future short rate. The *correction* is the convexity term $\frac{\sigma^2}{2a^2}(1-e^{-at})^2$: it is zero today, grows with the horizon and saturates, scales with σ^2 (vanishing entirely when volatility is switched off), and is the only place the volatility parameters touch the curve fit at all.

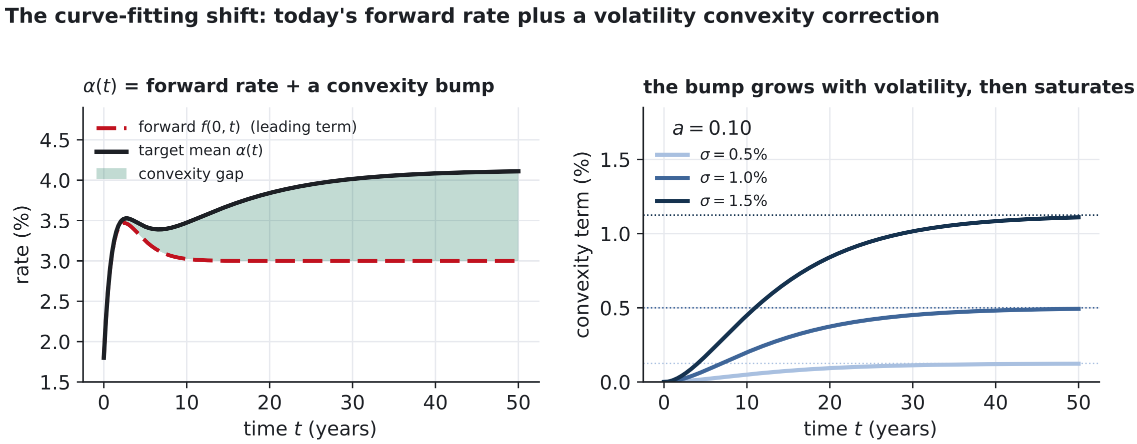

Figure 11: The structure of the curve-fitting shift $\alpha(t)$. *Left:* $\alpha(t)$ (black) is today's forward curve $f(0,t)$ (red dashed) plus a convexity bump (green gap). *Right:* the bump on its own—zero today, growing with horizon and saturating at $\sigma^2/2a^2$ (dotted), and scaling with σ^2, so more volatility means a bigger correction. (Drawn with $a = 0.10$ and an enlarged σ, over a long horizon, to make the gap and its saturation visible; at the $a = 0.12$, $\sigma = 1\%$ of Figure 8 it is the slim gap seen there between the mean and the forward.)

Why the convexity bump is there. The correction is the fingerprint of randomness, and one fact—Jensen's inequality—explains it exactly. Collect the randomness in the bond into a single number, the *average* short rate the model will experience over $[0,T]$,

$$\bar{r}_T \;=\; \frac{1}{T}\int_0^T r_s \,\mathrm{d}s,$$

which is random because the path is. The bond is then worth the average over all scenarios of its discount factor,

$$P(0,T) \;=\; \mathbb{E}\big[\,D(\bar{r}_T)\,\big], \qquad D(\bar{r}) = e^{-\bar{r}T},$$

and the object to look at is that discount factor $D(\bar{r})$. *It is the discount factor that is convex*: as a function of the rate, $D(\bar{r}) = e^{-\bar{r}T}$ bows upward, its second derivative $T^2e^{-\bar{r}T}$ being positive at every rate. That single property is what makes averaging and discounting fail to commute.

Jensen's inequality, in one line. For any *convex* function g and any random variable X,

$$\mathbb{E}[g(X)] \;\geq\; g\big(\mathbb{E}[X]\big)\;:$$

the average of the function lies *above* the function of the average. The reason is purely geometric—a convex curve sits below its chords (Figure 12). Take two equally likely outcomes for the average rate, one low and one high. The bond's value is the average of their two discount factors, which is the *midpoint of the chord* joining the two points on the curve; discounting instead at the mean rate gives the point on the curve directly below that midpoint. Because the curve bows beneath the chord, the true average sits higher: the bond is worth *more* than discounting at the expected rate would suggest. The intuition is the asymmetry of discounting—a rate that turns out high discounts the bond a little harder, but a rate that turns out equally low discounts it *far* more softly, and averaged together the soft side wins.

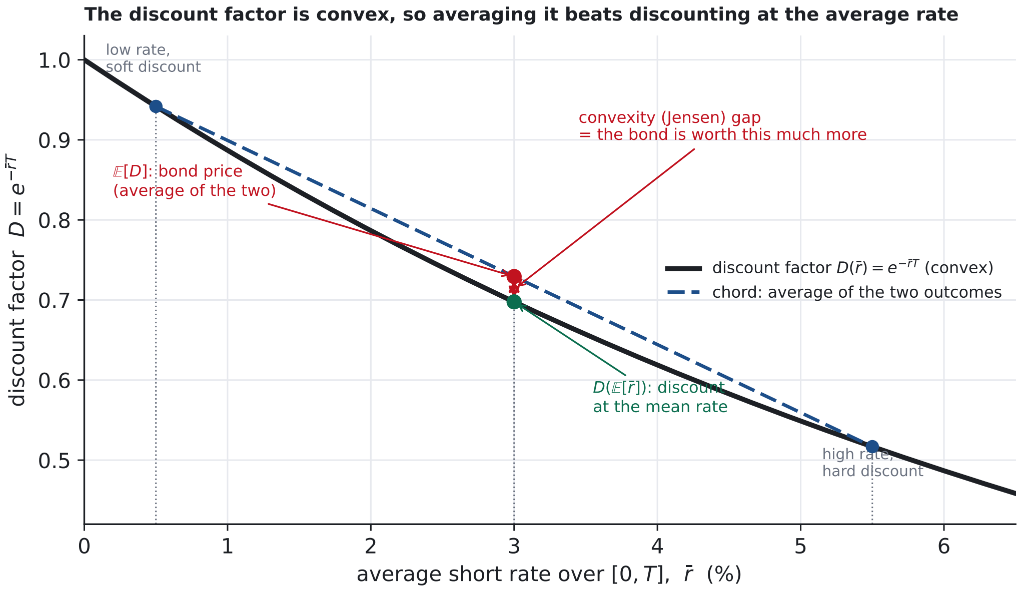

Figure 12: Jensen's inequality, which is the whole reason for the convexity bump. The discount factor $D(\bar{r}) = e^{-\bar{r}T}$ is a *convex* function of the average rate $\bar{r}$ (here $T = 12$, two equally likely outcomes at $\bar{r} = 0.5\%$ and 5.5%). Averaging the discount factor—the bond price $\mathbb{E}[D]$ (red), the midpoint of the chord—lands *above* discounting at the mean rate, $D(\mathbb{E}[\bar{r}])$ (green, on the curve). The gap is the convexity bonus that volatility hands the bondholder; the model must give it back to fit the curve.

Applied here, Jensen's inequality reads $\mathbb{E}\left[e^{-\bar{r}_T T}\right] > e^{-\mathbb{E}[\bar{r}_T]\,T}$: pure volatility makes a bond worth more than its expected-rate discount. If the model left its mean short rate at the bare forward $f(0,t)$, this bonus would push its bond prices *above* the market's—it would fit the curve too richly. To cancel the bonus and land back exactly on the market price, the model lifts its mean short rate slightly above the forward, by precisely the convexity term of (12): a higher mean rate discounts a little harder and undoes the extra value. The more volatility, the larger the Jensen bonus and the larger the lift needed—the σ^2 scaling seen in Figure 11 (right).

The exact-fit property. The crucial structural fact is that (12) holds for *any* choice of a and σ: whatever volatility parameters one picks, the shift $\alpha(t)$ recomputes itself to keep the curve repriced. The vanilla discount curve therefore does *not* pin down the volatility parameters—it is matched automatically—which is precisely why those parameters are free to be calibrated to something else (option prices), exactly as flagged in Section 3. Figure 13 confirms the property numerically: simulate the model, average the pathwise discount factors, and the recovered curve sits on top of the input curve to within Monte-Carlo noise.

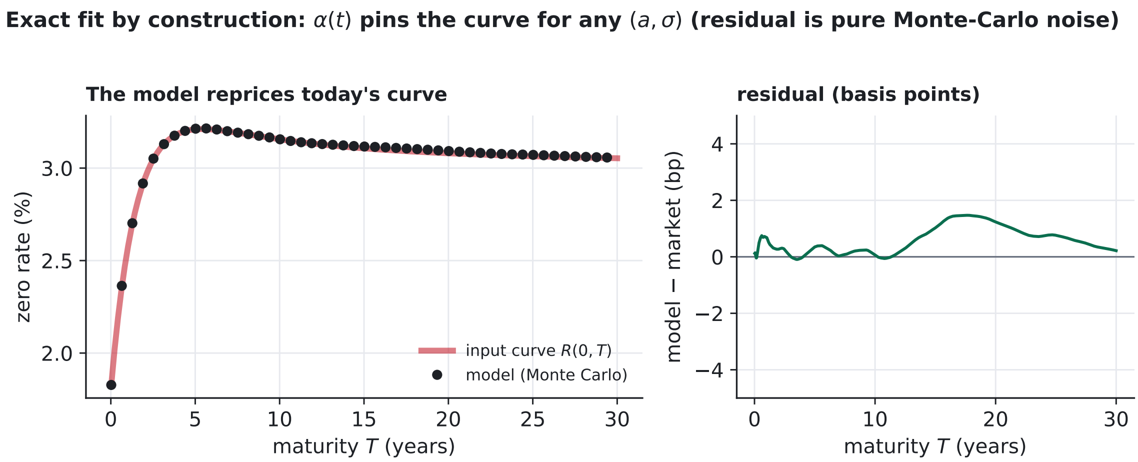

Figure 13: The exact-fit property, checked by Monte Carlo. *Left:* the zero curve recovered from the simulated model (dots) lies on the input curve (red). *Right:* the residual is a couple of basis points and structureless—pure Monte-Carlo noise, not model error. The shift $\alpha(t)$ guarantees the fit analytically, for any (a, σ).

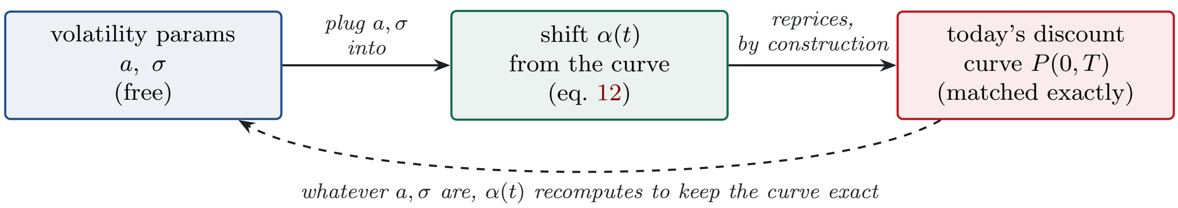

Figure 14: The exact layer. The deterministic shift $\alpha(t)$ is a slave to the curve: for any volatility parameters it is the unique function that reprices $P(0, T)$. So the curve never constrains a and σ—they are spent on the dynamics instead (Section 8).

The traditional way: fitting the drift $\theta(t)$

The same fit can be written directly in the original SDE (9), with the time-dependent drift $\theta(t)$ as the free object instead of the shift. Differentiating the decomposition $r = x + \alpha$ and matching drifts gives the relation $\theta(t) = \alpha'(t) + a\,\alpha(t)$, and substituting the shift (12) produces the classical Hull–White result

$$\theta(t) = \frac{\partial f(0, t)}{\partial t} + a\, f(0, t) + \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big). \tag{13}$$

This is the form in Hull and White's original paper and in many textbooks. It says exactly what (12) says—it reprices the curve for any (a, σ)—but far less legibly. It carries a derivative of the input forward curve, $\partial f(0, t)/\partial t$, so it is sensitive to how smoothly that curve was built, and it offers no clean "forward plus a small bump" reading. This is why almost no implementation forms $\theta(t)$ explicitly: one simulates the simple zero-mean x_t and adds the precomputed shift $\alpha(t)$. The drift form is worth recognising when you meet it in the literature, but the shift $\alpha(t)$ is the one to think with.

6 From the short rate to bond prices

Pricing anything beyond the curve itself needs bond prices at *future* dates, $P(t, T)$ for $t > 0$—the discount factor a future version of the world will quote. The reason the Hull–White model is so workable is that this comes out in closed form. The model is *affine*: the log bond price is linear in the short rate,

$$P(t,T) = A(t,T)\, e^{-B(t,T)\, r_t}, \qquad B(t,T) = \frac{1-e^{-a(T-t)}}{a}, \tag{14}$$

with a deterministic coefficient $A(t,T)$ fixed by today's curve,

$$A(t,T) = \frac{P(0,T)}{P(0,t)} \exp\Big(B(t,T)\, f(0,t) - \frac{\sigma^2}{4a}\big(1-e^{-2at}\big)\, B(t,T)^2\Big). \tag{15}$$

(Appendix A derives these from the bond pricing equation; they fall out of two simple ordinary differential equations.) The whole future curve is thus a deterministic function of the single number r_t: knowing today's curve and the current short rate, every discount factor follows. The coefficient $B(t,T)$ is the *sensitivity* of the log bond price to the short rate, and its shape—rising from 0 at $t = T$ toward $1/a$ for long maturities—is the model's term structure of rate sensitivity. It reappears everywhere below. Figure 15 plots the two coefficients against maturity.

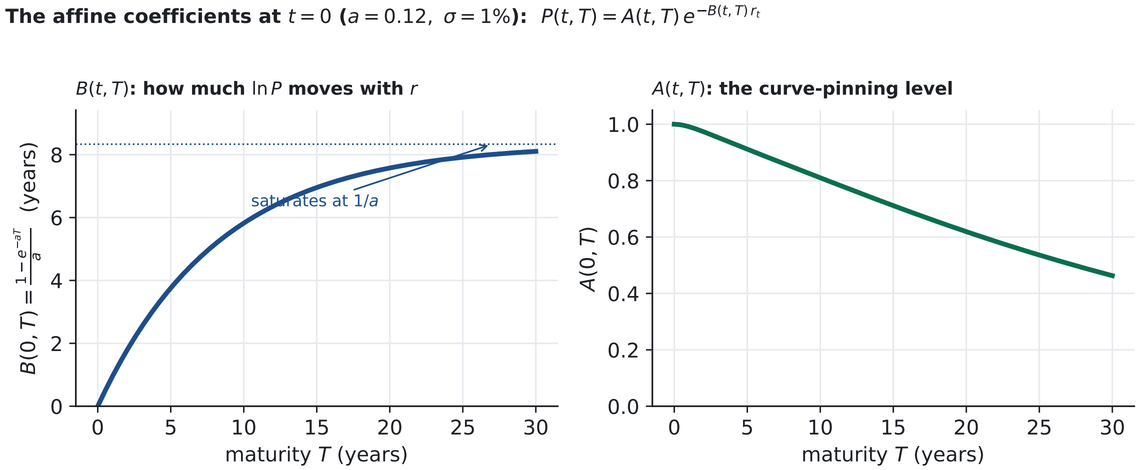

Figure 15: The two affine coefficients as functions of maturity (valuation time $t = 0$, base case $a = 0.12$, $\sigma = 1\%$). *Left:* $B(t,T)$ of (14), the sensitivity of the log bond price to the short rate, rises from 0 at $T = t$ and saturates at $1/a$—the model's term structure of rate sensitivity (here $1/a \approx 8.3$). *Right:* $A(t,T)$ of (15), the deterministic level coefficient set by today's curve, which falls off with maturity like the discount curve it encodes. Together they deliver every future bond price $P(t,T) = A(t,T)\, e^{-B(t,T) r_t}$ from the single number r_t. (At $t = 0$ the σ^2 term in A vanishes; for $t > 0$ it adds a small convexity correction.)

One number, one shape of curve move. The affine formula (14) has an immediate and far-reaching consequence. Because the whole curve is a deterministic function of the single number r_t, a change in r_t moves *every* maturity at once—and always by the same shape. Bump the short rate by a small amount and the forward rate at maturity T shifts by a factor $e^{-a(T-t)}$: fully at the front, fading to almost nothing at the long end (Figure 16, left). And that profile is *fixed*—every possible move of the curve is some scalar multiple of the one shape $e^{-a(T-t)}$ (Figure 16, right). A one-factor model can raise, lower and tilt the curve through this single mode, but it cannot bend the front and the back in opposite directions; that rigidity is the one-factor limitation we cash out in Section 11. For now it has a useful upside: it makes the curve's volatility structure completely explicit.

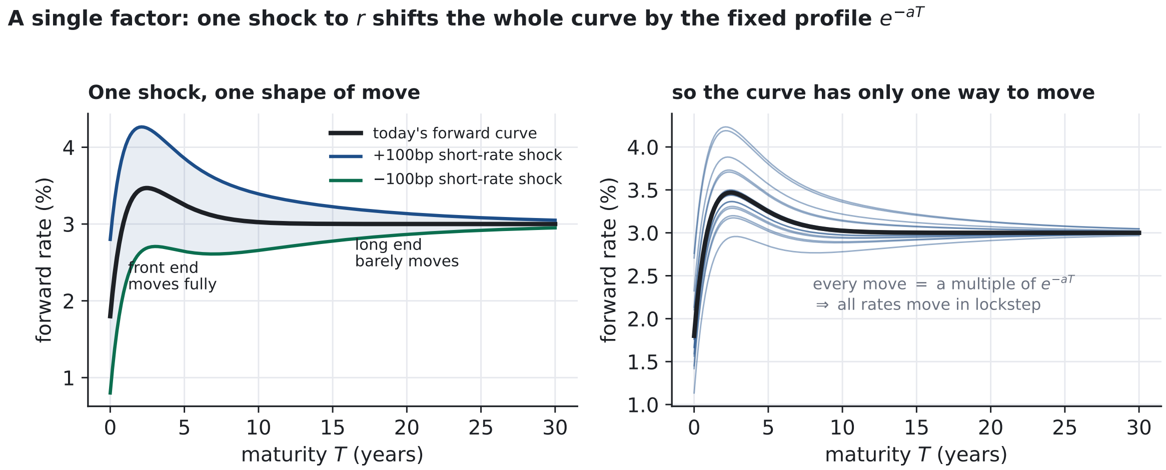

Figure 16: One factor, one shape of move. *Left:* a ±100bp shock to the short rate shifts the whole forward curve by the profile e^{-aT}—the front end moves fully, the long end barely. *Right:* every possible curve move is a scalar multiple of that single profile, so the curve has just one way to move and all rates move in lockstep—the seed of the one-factor limitation (Section 11). Here $a = 0.10$.

Mean reversion shapes the volatility of the curve. "Differentiating the affine form" means taking $-\partial_T \ln P(t,T)$ of (14) to recover the instantaneous forward rate. Since $\ln P(t,T) = \ln A(t,T) - B(t,T)\, r_t$,

$$\begin{aligned} f(t,T) &= -\partial_T \ln P(t,T) \\ &= \underbrace{-\,\partial_T \ln A(t,T)}_{\text{deterministic}} + \big(\partial_T B(t,T)\big)\, r_t. \end{aligned} \tag{16}$$

Since $\partial_T B(t,T) = e^{-a(T-t)}$, the short rate enters the T-maturity forward rate with the coefficient $e^{-a(T-t)}$. The only randomness is the Gaussian r_t, whose instantaneous volatility is σ; hence the instantaneous volatility of the T-maturity forward rate seen from time t is $\sigma\, e^{-a(T-t)}$. The same shock σ to the short rate thus moves a near forward rate fully but a distant one hardly at all, damped by the mean reversion (Figure 17, left). This is the single most important qualitative consequence of a. A small a means shocks propagate far out the curve, so long rates are nearly as volatile as short ones; a large a means shocks die quickly, so only the front end moves and the long end is calm. Equivalently (Figure 17, right), a caps the stationary spread of the short rate itself, via (11). The mean-reversion speed is, in effect, the model's *term structure of volatility* knob.

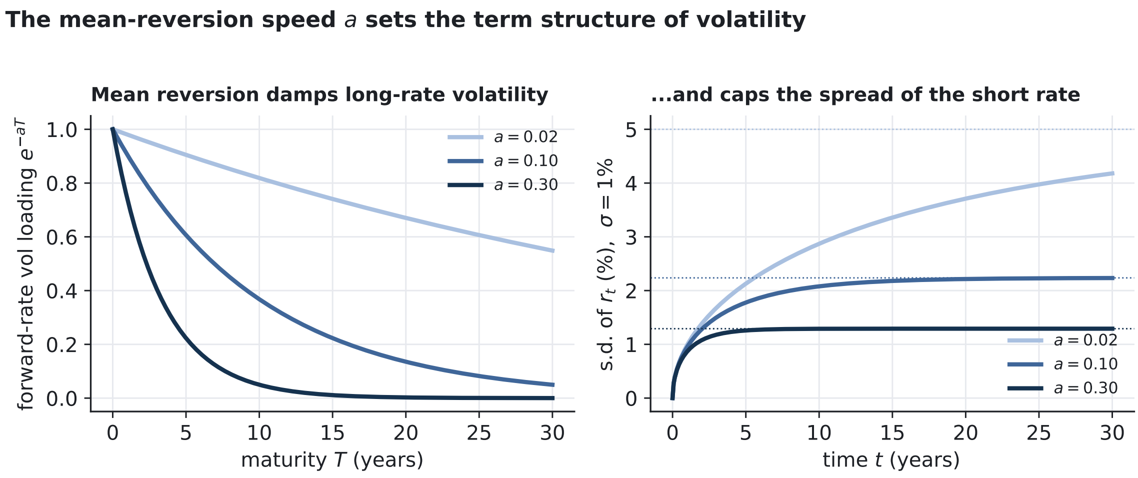

Figure 17: What the mean-reversion speed a controls. *Left:* the volatility of the T-maturity forward rate is $\sigma\, e^{-aT}$; a larger a damps distant rates more, steepening the decay. *Right:* the standard deviation of the short rate saturates at $\sigma/\sqrt{2a}$ (dotted), reached sooner and lower for larger a. One parameter sets how volatility is distributed across maturities *and* how wide the long-run rate distribution is.

7 What the one-factor model prices: the vanillas

The affine structure makes a surprising amount analytic. We build the vanilla interest-rate options from the ground up, in the order they nest: a caplet is an option on one zero-coupon bond, a cap is a strip of those, and a European swaption is an option on a coupon bond that Jamshidian splits back into single-bond options. So the whole vanilla book rests on one closed form—a European option on a zero-coupon bond—and the rest is bookkeeping. We take our time over that one engine, because its volatility input is the single number everything else inherits.

Caplets, and why they are bond options. Begin with the simplest interest-rate option, the *caplet*: it is a *call on a single forward interest rate*—one period of a cap. It involves three dates, worth keeping straight (Figure 18): *today*, $t = 0$, when we price it; the *expiry* T_O, at which the relevant rate is observed and the option pays off or expires; and the end of the accrual period, T_B, at which the cash actually changes hands. The subscripts are a memory aid—O for the *option*'s expiry, B for the *bond* that will turn out to underlie it. Write $\tau = T_B - T_O$ for the accrual length (a quarter or half-year, usually) and F for the simply-compounded forward rate over $[T_O, T_B]$, the rate fixed at T_O. The caplet with strike K pays

$$\tau\,(F - K)^+ \quad \text{at } T_B\,:$$

it caps borrowing at K, paying the excess when the rate fixes above K and nothing when it fixes below.

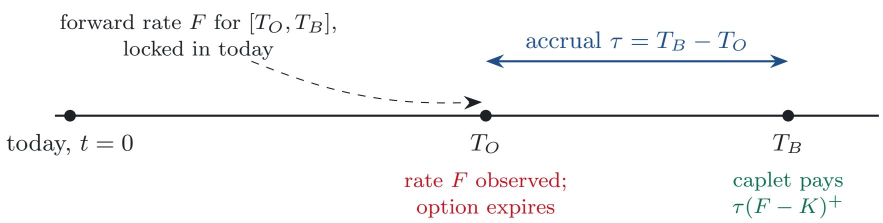

Figure 18: The three dates of a caplet. The forward rate F for the period $[T_O, T_B]$ is known today; it is *observed* (fixed) at the expiry T_O, and the resulting payoff $\tau(F - K)^+$ is paid at the end of the accrual period T_B. Pricing happens at $t = 0$. The subscripts read Option-expiry and Bond-maturity.

Why is this—a call on a *rate*—also an option on a *bond*? Because a high rate and a cheap bond are the same event, and the link is just the definition of the simply-compounded rate. Lending \$1 over the accrual period at the rate F returns $1 + F\tau$ at T_B; the same loan made through the bond market—buy $1/P(T_O, T_B)$ of the T_B-bond, which costs \$1 at T_O and pays $1/P(T_O, T_B)$ at T_B—must return the same amount, so

$$1 + F\tau \;=\; \frac{1}{P(T_O, T_B)}.$$

A large F therefore means a small $P(T_O, T_B)$: rates up, bond down. A caplet, which pays when F is *high*, thus pays when the bond is *cheap*—which is precisely a *put* on the bond. To make that exact, value the T_B payoff back at T_O by multiplying by $P(T_O, T_B)$, and substitute $1 + F\tau = 1/P(T_O, T_B)$:

$$\begin{aligned}
\underbrace{P(T_O, T_B)\,\tau\,(F - K)^+}_{\text{caplet payoff, valued at } T_O} &= \big(1 - (1 + K\tau)\,P(T_O, T_B)\big)^+ \\
&= (1 + K\tau)\Big(\tfrac{1}{1+K\tau} - P(T_O, T_B)\Big)^+.
\end{aligned}$$

The last line is $(1 + K\tau)$ identical *puts* on the zero-coupon bond $P(\cdot, T_B)$, each expiring at T_O and struck at $1/(1 + K\tau)$. So a caplet is, with no approximation, a small basket of bond puts—a call on the rate is a put on the bond—and pricing one bond put prices the caplet, and, summed over a schedule, the whole cap. Everything now reduces to one task: pricing a European option on a zero-coupon bond.

What "bond-option volatility" is, and why we need it. An option is worth more when its underlying is more uncertain—the one fact behind every option formula. To price the bond put, then, we must know how uncertain the bond's price is by the time the option expires: how widely $P(T_O, T_B)$ is spread across the ways the world might unfold between now and T_O. That single number—the standard deviation of the (log) bond price accumulated to expiry—is the *bond-option volatility*, written Σ_P. It is the exact counterpart of the $\sigma\sqrt{T}$ in Black–Scholes: a bond that might end up far below the strike or far above it (large Σ_P) carries valuable options, while a bond whose price is nearly pinned down by T_O (small Σ_P) carries nearly worthless ones. Figure 19 is the whole idea in one picture: the bond's price fans out as expiry approaches, and the width of that fan at T_O is what the option prices off. So the real task is to find Σ_P.

Bond-option volatility Σ_P: how wide the bond's price is when the option expires

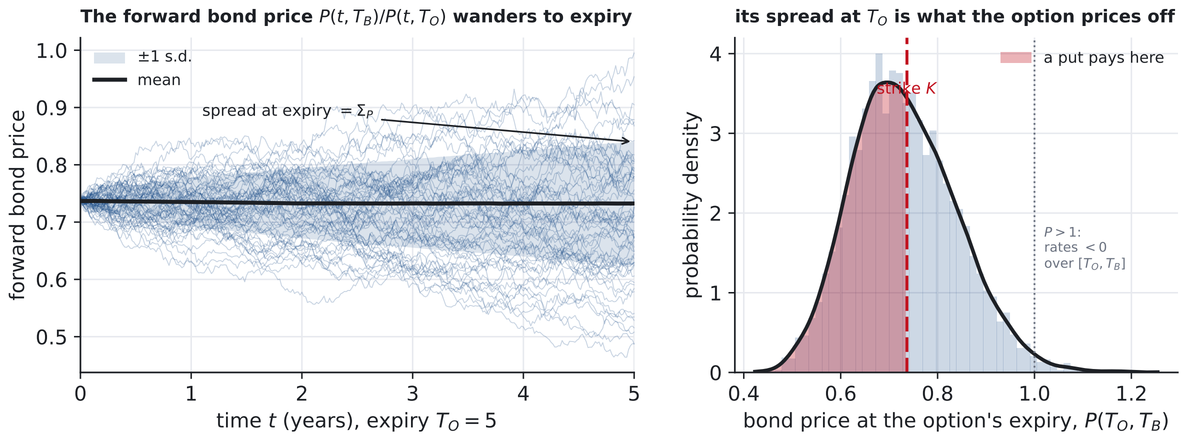

Figure 19: What the bond-option volatility Σ_P measures. *Left:* the forward price of the underlying bond, $P(t, T_B)/P(t, T_O)$, wanders as the option's expiry T_O approaches; the spread it has built up by T_O is Σ_P. *Right:* at T_O that spread *is* the distribution of the bond price $P(T_O, T_B)$, and the bond put pays on the part of it below the strike K (shaded). A wider distribution—a larger Σ_P—means a more valuable option. Note that $P(T_O, T_B)$ is the bond's price at the *option's* expiry T_O, with the tenor $T_B - T_O$ still to run—not at the bond's own maturity T_B, where it pays par—so it can exceed 1: that right tail is exactly the scenarios where rates over the remaining tenor turn negative, which the Gaussian model permits. (Drawn for a 5-year option on the 10-year-forward bond, so the spread is clearly visible.)

The forward bond price, and the forward measure. One refinement before the arithmetic. The option settles at T_O, so the quantity that matters is the bond's price *as seen from expiry*, which is cleanest to track through the *forward bond price*—the price one could lock in today for delivery of the T_B-bond at T_O,

$$G(t) \;=\; \frac{P(t, T_B)}{P(t, T_O)}.$$

At expiry the front bond matures at par, $P(T_O, T_O) = 1$, so $G(T_O) = P(T_O, T_B)$: the forward price simply becomes the bond price the option pays off against. Working with G unlocks a standard and powerful device, the *forward measure.* If we quote every value in units of the T_O-bond—using $P(\cdot, T_O)$ as the unit of account, or *numeraire*, in place of cash—then G has *no drift*: it is a fair bet, a martingale. Intuitively, measuring everything in T_O-bonds strips out the deterministic pull of discounting and leaves only the randomness, which is all an option's value depends on. Under this measure the bond option becomes a textbook *Black* calculation—a European option on a lognormal forward—needing only that forward, G, and its total volatility, Σ_P. The change-of-numeraire mechanics and the Black formula are set out in Appendix C; here we need only supply the one input it asks for, Σ_P.

Computing Σ_P. We get there in three short steps: the volatility of a single bond, then of the forward bond, then its accumulation to expiry.

Step 1: the volatility of one bond. The affine form (14) hands it to us almost for free. Take logs, $\ln P(t, T) = \ln A(t, T) - B(t, T)\, r_t$, whose only random term is $-B(t, T)\, r_t$. A unit move in the short rate moves the log bond price by $-B(t, T)$, and since the short rate has volatility σ, the T-bond's log price has instantaneous volatility $\sigma\, B(t, T)$. That is just the rate volatility σ *geared up* by the bond's sensitivity $B(t, T)$—longer bonds, with larger B, move more.

Step 2: the volatility of the forward bond. The option's underlying is the forward bond price $G = P(t, T_B)/P(t, T_O)$. Being a ratio, its log is the *difference* of the two bonds' log prices, so it carries the difference of their volatilities, $\sigma\big(B(t, T_B) - B(t, T_O)\big)$—and that difference collapses neatly:

$$
\begin{aligned}
B(t,T_B) - B(t,T_O) &= \frac{e^{-a(T_O-t)} - e^{-a(T_B-t)}}{a} \\
&= e^{-a(T_O-t)}\, B(T_O,T_B).
\end{aligned}
$$

So at time t the forward bond price has instantaneous volatility $\sigma\, e^{-a(T_O-t)}\, B(T_O,T_B)$. Read it as two pieces: the *size* $B(T_O,T_B)$, the sensitivity of the forward bond to the rate; and the *factor* $e^{-a(T_O-t)}$, which says that shocks arriving long before expiry have largely decayed by the time the option pays—mean reversion again.

Step 3: accumulate to expiry. It is *variance*, not volatility, that adds up over time: the total variance to expiry is the instantaneous variance integrated over the option's life, and Σ_P is its square root. So we square the step-2 volatility, integrate, and take the root:

$$
\begin{aligned}
\Sigma_P^2 &= \int_0^{T_O} \left(\sigma\, e^{-a(T_O-t)}\, B(T_O,T_B)\right)^2 \mathrm{d}t \\
&= \sigma^2\, B(T_O,T_B)^2\, \frac{1-e^{-2aT_O}}{2a},
\end{aligned}
$$

which gives

$$
\Sigma_P \;=\; \sigma \;\; \underbrace{\sqrt{\tfrac{1-e^{-2aT_O}}{2a}}}_{\text{accumulated to } T_O} \;\; \underbrace{B(T_O,T_B)}_{\text{bond sensitivity}} \;. \tag{17}
$$

So Σ_P is a product of three plain quantities: the short-rate volatility σ; the sensitivity $B(T_O,T_B)$ that turns rate moves into bond-price moves; and an accumulation factor that grows with the option's life but is capped by mean reversion. (It is the same accumulated-variance integral that returns, with a time-dependent $\sigma(t)$, in the calibration bootstrap of §8.)

The bond option, priced. With Σ_P in hand, Black's formula (Appendix C) prices the option directly. The option to buy the T_B-bond at expiry T_O, struck at K, is worth

$$
P(0,T_B)\,\Phi(d_+) - K\,P(0,T_O)\,\Phi(d_-), \qquad d_\pm = \frac{\ln\!\big(P(0,T_B)/(K\,P(0,T_O))\big)}{\Sigma_P} \pm \tfrac{1}{2}\Sigma_P,
$$

with the put following by put–call parity. That is the engine: the caplet used it above, the swaption will use it below, and the worked example next puts numbers to it.

Worked example: an at-the-money caplet. Price the caplet over $[T_O,T_B] = [2,\ 2.5]$ years—a six-month rate, two years forward—on this note's curve, with $a = 0.05$ and $\sigma = 1\%$. Read off two discount factors and follow (17):

- discount factors $P(0,2) = 0.9429,\ \ P(0,2.5) = 0.9267$;
- ATM forward rate $F = \big(P(0,2)/P(0,2.5) - 1\big)/0.5 = 3.49\%$;
- bond-option sensitivity $B(T_O,T_B) = (1 - e^{-0.05\times 0.5})/0.05 = 0.4938$;
- accumulated log-bond volatility $\Sigma_P = \sigma\sqrt{(1-e^{-2\times0.05\times2})/(2\times0.05)}\,B(T_O,T_B) = 0.00665$;
- the caplet is $(1+F\tau)$ puts on the 2.5-year bond struck at $K = 1/(1+F\tau) = 0.9828$; the Black bond-put formula at volatility Σ_P returns a put worth 0.002458, so the caplet is worth $(1+F\tau)\times 0.002458 = 0.00250$—about 25 basis points of notional;
- quoted as an implied volatility that is a Black (lognormal) vol of 27.6%, or equivalently a normal (Bachelier) vol of 96 bp.

Every line is arithmetic; no simulation is involved. The same recipe, summed over a strip, prices caps and floors, and via Appendix E European swaptions.

Caps and floors, by strip. A *cap* insures a floating-rate borrower: over a schedule of reset periods it pays out, in each period, whenever the floating rate fixes above the strike K. So it is nothing more than a *strip of caplets*—one for each accrual period—and a *floor*, which protects a lender against rates falling below K, is the matching strip of *floorlets* (bond *calls*). Each piece prices by the bond-option formula above, so the cap (resp. floor) is simply their sum. Two checks come for free. A floorlet is the put–call partner of its caplet on the same bond, so caplet − floorlet is the discounted forward payoff $\tau(F - K)$; summed over the schedule, *cap* − *floor* = *swap*—the model's options obey put–call parity against the underlying swap automatically. The whole cap/floor market is thus priced analytically off one Σ_P per period, and the resulting at-the-money caplet volatilities (Figure 20) show the model's signature: with a single constant σ, the *shape* of the volatility term structure across expiries is set almost entirely by the mean reversion a.

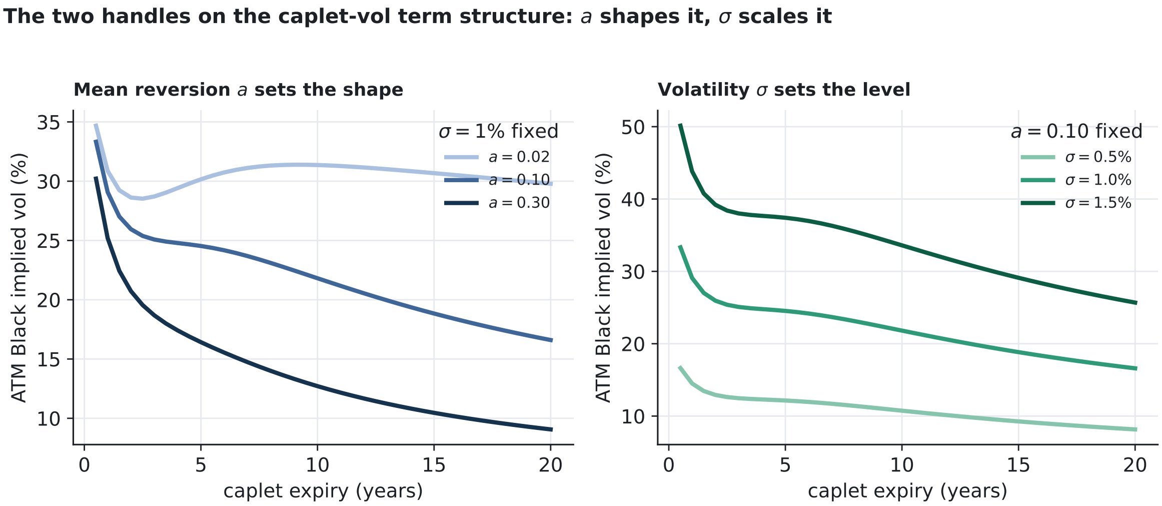

Figure 20: The two handles a one-factor model has on the at-the-money caplet (Black) volatility term structure. *Left:* varying the mean reversion a at fixed $\sigma = 1\%$ changes the *shape*—faster mean reversion pulls the curve down and steepens its decline with expiry. *Right:* varying the volatility σ at fixed $a = 0.10$ changes only the *level*, scaling the whole curve up or down while leaving its shape essentially unchanged. So one constant pair (a, σ) buys a level and a shape and no more—which is why matching a whole market term structure needs the time-dependent $\sigma(t)$ of §8, and a genuine hump needs the second factor of Part II.

So what is that term structure *for*? It is the market's own quoting language, and it is why this picture is really about calibration. Caps and caplets are quoted not as prices but as a curve of at-the-money implied volatilities—one number per expiry—so pricing any cap means reading values off that curve. The model must therefore *reproduce* the market's quoted curve, and (17) lays bare the only handles it has for the job: an overall *level*, set by σ, and a *decay with expiry*, set by a. Two constants buy only a two-parameter family of shapes—enough to match a level and a slope, but not a whole curve that bends or humps. Closing that gap is precisely the task of calibration (§8), where σ is allowed to vary in time, $\sigma(t)$, and bootstrapped expiry by expiry so the model's curve passes through every quote; a genuinely *humped* term structure, which no single a can bend to, is one of the reasons for the second factor of Part II.

Swaps, in bond prices. A swaption is an option on a *swap*, so we first need the swap itself in bond-price terms. An interest-rate swap exchanges, over a schedule of dates $T_1 < \cdots < T_n$ with accruals τ_i, a stream of *fixed* coupons at a rate K for the *floating* rate that fixes each period—the standard way to turn a floating exposure into a fixed one, or back. Appendix D sets out the basic mechanics of a vanilla swap in detail—the payment schedule, the fix-in-advance, pay-in-arrears timing, and the payer/receiver and settlement conventions; here we need only its value in bond prices. Both legs are bonds in disguise. The *fixed leg* is a *coupon bond*—the fixed coupon on each date, plus the notional returned at T_n—worth the discounted sum of its cashflows,

$$C(t) = \sum_i c_i \, P(t, T_i),$$

where $c_i = K\tau_i$ is the coupon at T_i and the final cashflow also returns the notional (taken as 1). The *floating leg* telescopes, by no-arbitrage, to $P(t, T_0) - P(t, T_n)$—receiving the floating rate is exactly rolling a deposit—which for a spot-starting swap ($T_0 = t$) is $1 - P(t, T_n)$. A *payer* swap (pay fixed, receive floating) is therefore worth the difference of the two legs,

$$\begin{aligned} V_{\text{pay}}(t) &= \big(1 - P(t, T_n)\big) - K \textstyle\sum_i \tau_i \, P(t, T_i) \\ &= 1 - C(t), \end{aligned}$$

which gains when rates rise and the fixed-leg bond $C(t)$ falls below par. The fixed rate that makes the swap worth nothing today—the *forward swap rate*—is

$$S(t) = \frac{P(t, T_0) - P(t, T_n)}{\sum_i \tau_i \, P(t, T_i)}, \tag{18}$$

the floating leg over the *annuity* $\sum_i \tau_i P(t, T_i)$ (the present value of \$1 of fixed coupon). A swap, or a swaption, is *at the money* when $K = S(t)$—the swap analogue of a caplet struck at the forward rate.

European swaptions, by Jamshidian. A *payer swaption* is the option to enter that payer swap at the start date T_O, so it pays $\big(1 - C(T_O)\big)^+$ at expiry—a *put* on the fixed-leg coupon bond C, struck at par. This too has an exact decomposition, due to Jamshidian [5], resting on the one-factor luxury we have leaned on throughout: because every bond price $P(T_O, T_i)$ is a monotone function of the single state r_{T_O}, the coupon bond $C(T_O) = \sum_i c_i \, P(T_O, T_i)$ is monotone in r_{T_O} too, so it crosses par at exactly one *critical rate* r^*, the rate at which $C(T_O) = 1$. Set each bond's strike to the value it takes there, $K_i = P(T_O, T_i)\big|_{r_{T_O} = r^*}$. Because every bond crosses its own K_i at that *same* rate r^*, the put on the coupon bond splits exactly into a portfolio of puts on the individual zero-coupon bonds:

$$\big(1 - C(T_O)\big)^+ \;=\; \sum_i c_i \, \big(K_i - P(T_O, T_i)\big)^+. \tag{19}$$

Each piece is a European put on a zero-coupon bond, priced in closed form by the bond-option formula at its own volatility (17); the matching receiver swaption is the call dual. Appendix E gives the general decomposition and its derivation. So the entire European swaption grid—the main calibration instrument—is analytic. The geometry is worth holding onto (Figure 21): one state variable, one threshold. That single threshold is exactly the luxury two factors break (§14).

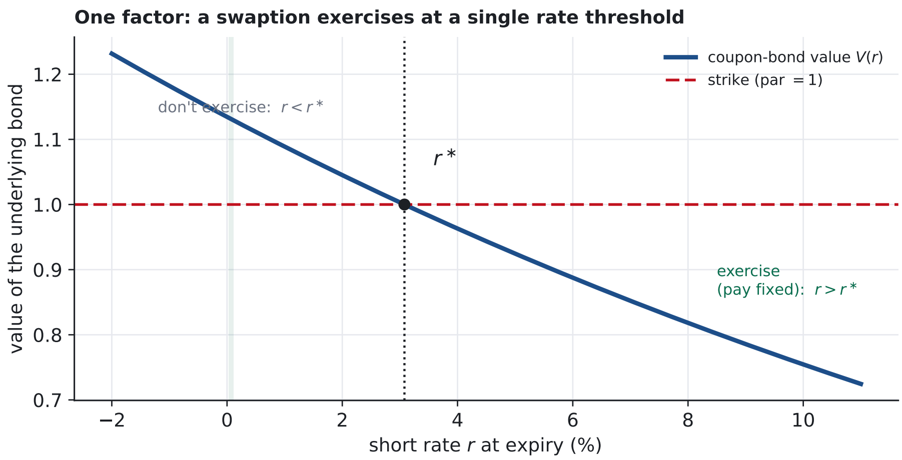

Figure 21: Why one factor makes swaptions analytic. The value of the underlying coupon bond is monotone in the single state r_{T_O}, so it crosses the strike at exactly one rate r^*—exercise is decided by whether the rate is above or below a single threshold. Jamshidian's decomposition (Appendix E) then rewrites the swaption as a portfolio of bond options struck at that one rate. With two factors this single crossing point swells into a whole boundary curve (Section 14), and the decomposition breaks.

8 Calibrating the one-factor model

Calibration has the same two layers as the model. The *inner, exact layer* is $\alpha(t)$: given any (a, σ) it is computed once from the curve by (12), and the curve is matched exactly—no optimisation, no error. The *outer layer* is the genuine fit: choose the dynamics parameters a and σ so the model's *option* prices match the market's. Everything in this section is about that outer layer, and about doing it with discipline. The good news, which we lean on throughout, is that the prices being matched are the closed forms derived in §7—the bond-option price and, through Jamshidian's decomposition (Appendix E), the European swaption grid—so the fit is a fast analytic least-squares problem with no Monte-Carlo noise to fight.

The market data: a grid of volatilities. The raw material is a grid of quoted option volatilities—always *at-the-money* options, struck at the prevailing forward rate, not a cross-section of strikes. (The smile across strikes is the separate topic of §9.) For swaptions it is a rectangular table indexed by *option expiry* down one side (1m, 3m, 6m, 1y, 2y, . . . , out to 30y) and *underlying swap tenor* across the top (1y, 2y, 5y, 10y, 30y): each cell is the implied volatility of the at-the-money European swaption with that expiry and tenor—a "5y into 10y" is the option, five years out, to enter a ten-year swap. *At the money* means struck at the forward swap rate, the rate that makes the swap worth zero today, so the quote carries pure volatility and no intrinsic value. These are quoted today as *normal* (Bachelier) volatilities, in basis points of rate per root-year, rather than the *Black* (lognormal) volatilities of the previous era—a switch whose meaning for the Black pricing of §7 is worth a paragraph of its own (below).

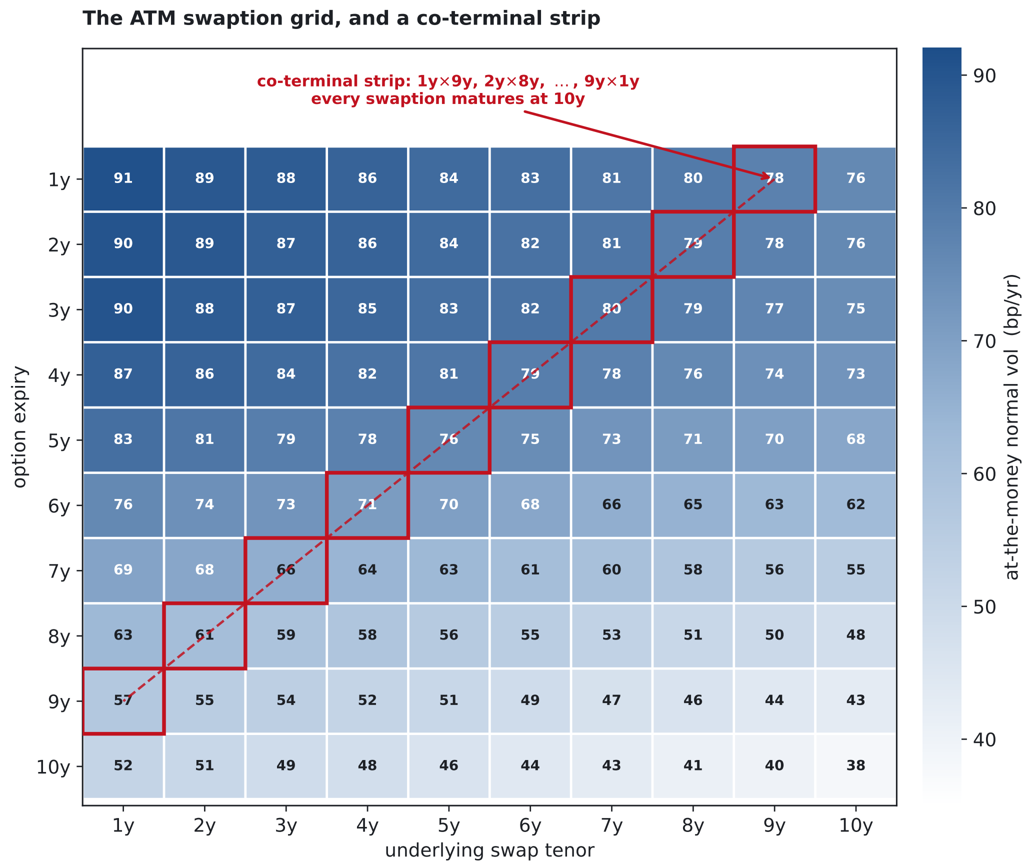

Figure 22: The at-the-money European swaption grid that anchors calibration (illustrative normal volatilities in basis points per year, on a schematic integer-year grid; not a live quote). Rows are option expiry, columns the tenor of the underlying swap, and each cell is the implied volatility of the at-the-money swaption with that expiry and tenor. One rarely fits the whole rectangle; the usual target is a one-dimensional strip cut from it. The outlined anti-diagonal is a *co-terminal* strip—"1y into 9y", "2y into 8y", . . . , "9y into 1y"—on which every swaption matures on the *same* date, ten years out; this is the strip a Bermudan swaption hedges with (discussed below). A single row (one expiry, varying tenor) or a single column (one tenor) is the other common choice.

Black, normal, and what is actually being matched. There is an apparent tension to clear up. In §7 we went to some lengths to price the vanillas with *Black*'s lognormal formula, and here the market quotes *normal* vols—so was that work spent on an obsolete convention? No, for two reasons. First, the model never produces a volatility at all; it produces a *price*, in currency, from the closed forms of §7. An implied volatility is only a quoting convention: an invertible map between that price and a single number through a reference formula. Black inverts the price through a lognormal reference, Bachelier through a normal one—two units for the same money, like quoting a bond by price or by yield. Calibration matches *prices*, so the market's choice of unit is a wrapper around the fit, not its substance; one converts each quoted normal vol back to a price through Bachelier and matches the model's price to it. Second, and more fundamentally, the Black formula we built prices an option on a *bond*, and a bond price is positive even when rates are negative—so that machinery never relied on positive rates. What broke when rates went through zero was only the separate habit of quoting an option on a *rate* as a lognormal volatility *of* that rate, a percentage of something that had turned negative. The model still prices through the always-positive lognormal bond; the market simply reports the answer in normal-vol units—which, by a convenient alignment, is the unit in which Hull–White's own smile sits closest to flat (§9). Where forwards are comfortably positive the lognormal quote still survives, often as a *shifted* (displaced) Black volatility on $F + s$ for a fixed shift s chosen to keep the reference forward positive.

Which slice of the grid, and the alternative. One rarely fits the whole rectangle; the usual swaption target is a one-dimensional strip cut from it. The *co-terminal* (or *diagonal*) strip holds the final maturity fixed and walks the expiry out—"1y into 9y, 2y into 8y, 3y into 7y, ..."—so every swaption in the strip matures on the same terminal date. The alternative target is *caps and floors*: a cap is a strip of caplets, each a little option on a single forward rate, and the at-the-money caplet volatilities form a strip indexed by expiry alone. Caps are simple and very liquid at the short end; swaptions reach further out and carry the coupon-bond optionality most exotic rate books actually hold. Which one you pick is not a matter of taste, and that is the next point.

Calibrate to the instruments your product hedges with. The governing principle is that you mark and hedge an exotic with the vanilla instruments it most resembles, so you calibrate to *those* instruments and let the model reprice them exactly. A Bermudan swaption is the canonical example. A Bermudan into a fixed swap can be exercised on a sequence of dates, and on each date it turns into the European swaption that runs from that date to the common final maturity: precisely the co-terminal strip. So one calibrates the Bermudan to its co-terminal Europeans (priced in §10), because those are the instruments the desk will hedge with and the ones whose mispricing would show up directly in the exotic [7]. A cap book, by the same logic, calibrates to the caplet strip; a product sensitive to one 10y rate ten years forward calibrates to the swaptions around that cell. The aim is not to fit *all* of the market—one factor cannot—but to fit exactly the part of it your risk lives in. Choosing the wrong strip is a more common error than any optimiser bug.

The two parameters separate by role. With the target chosen, the fit is well behaved because a and σ do different jobs:

- σ sets the overall *level* of volatility—how big the rate shocks are, hence the overall height of the caplet/swaption vols.
- a sets the *shape*: through the loading $\sigma e^{-a(T-t)}$ it controls how fast volatility decays with maturity (Figure 17), and it governs the correlation between rates of different maturities (§11). It tunes the *slope* of the volatility term structure—the relative pricing of options on long versus short rates.

But two constants can match only two features—a level and a slope—and a real market term structure is not so obliging. The standard refinement is to make the volatility *piecewise-constant in time*, $\sigma(t)$, and fit the whole at-the-money strip exactly, while keeping a a single global constant. This is the rates analogue of the two-layer split itself: a flexible time-dependent volatility absorbs the term structure, exactly as the flexible time-dependent drift absorbs the curve, while one hard-to-see global parameter carries the correlation.

The $\sigma(t)$ bootstrap, in detail. The reason $\sigma(t)$ can be fitted exactly, one quote at a time, is a sequential structure hiding in the closed forms. Order the option expiries $T_1 < T_2 < \cdots < T_n$ of the chosen strip and let $\sigma(t)$ be constant on each interval, equal to σ_i on $(T_{i-1}, T_i]$. The price of the option expiring at T_i depends on the volatility only through the variance the state accumulates up to that expiry,

$$v(T_i) \;=\; \int_0^{T_i} \sigma(u)^2\, e^{-2a(T_i-u)}\, \mathrm{d}u, \tag{20}$$

an integral that stops at T_i and runs no further—so the price at T_i sees $\sigma_1, \ldots, \sigma_i$ but is blind to every later piece. That makes the fit a clean forward sweep:

- fix the global a first (next paragraph);
- the first quote, at T_1, depends only on σ_1; invert the closed-form price for the single σ_1 that reproduces it;
- with σ_1 frozen, the quote at T_2 depends only on the new piece σ_2; solve for it;
- continue outward, each expiry T_i pinning exactly one new piece σ_i, until the strip is exhausted (Figure 23).

Each step is a one-dimensional root-find against an *analytic* price—no optimiser, no Monte Carlo, no noise. The strip is matched to machine precision, and because the recursion never revisits an earlier piece it is stable: a quote at the long end cannot disturb the short-end fit. This is the rates version of bootstrapping a discount curve, run on volatilities instead of discount factors.

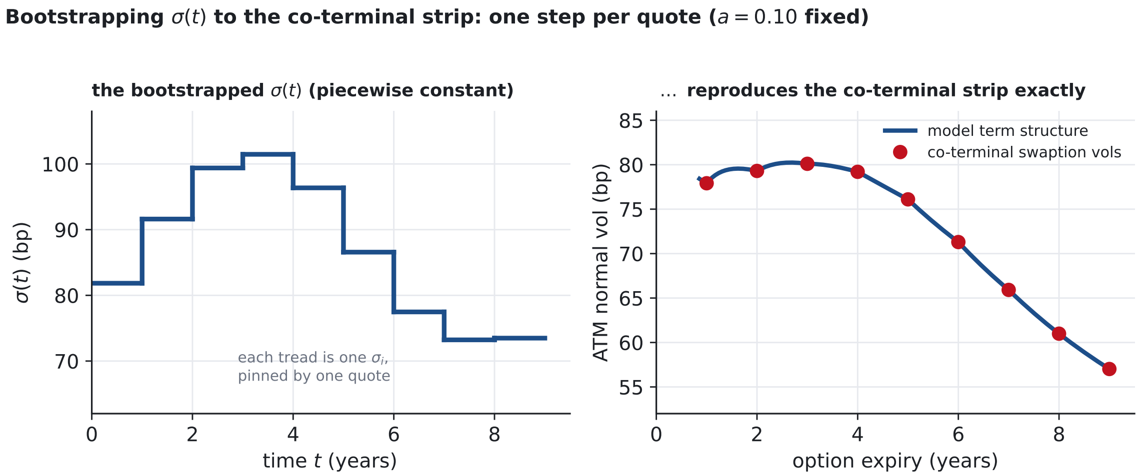

Figure 23: The $\sigma(t)$ bootstrap (illustrative, $a = 0.10$ fixed), run on the *co-terminal* strip outlined in Figure 22. *Right:* those same at-the-money normal volatilities (red), one quote per option expiry, $1y, 2y, \dots, 9y$, showing the characteristic *hump*—rising to a peak near three years, then declining. *Left:* the piecewise-constant $\sigma(t)$ recovered by the forward sweep—each tread σ_i is pinned by exactly one quote, since the price at expiry T_i sees the accumulated variance (20) only up to T_i. The model term structure (blue, right) passes through every market point exactly. Note that $\sigma(t)$ and the quoted strip need not share a shape: the recovered $\sigma(t)$ humps *higher, and a little later*, than the quoted strip, because each quote integrates the mean-reversion-weighted history $e^{-2a(T_i-u)}$—which discounts the early, decayed contributions—rather than the instantaneous level. A *sharper* market hump would force $\sigma(t)$ into an ever more contorted shape, one of the strains that motivates the second factor of Part II.

Pinning a properly. The mean reversion is the genuinely hard parameter. Against a single at-the-money strip a is *barely identified*: whatever slope a would produce, the bootstrapped $\sigma(t)$ can mimic by leaning its own time profile the other way, so the strip alone does not separate the two. What is left for a to say is everything $\sigma(t)$ *cannot* fake. Four reliable handles:

- *Relative pricing across tenors.* A single expiry strip pins a weakly, but the *full* grid does not: how long-tenor swaptions price relative to short-tenor ones at a given expiry is a slope $\sigma(t)$ cannot reproduce, because $\sigma(t)$ is a function of calendar time, not of tenor. Fitting across tenor as well as expiry is where a earns its keep.
- *The correlation it implies.* a sets the decorrelation of rates of different maturities (§11); anything whose value is a correlation speaks directly to it.
- *Historical estimation.* The speed at which a shock decays out of the curve is visible in the time series of rates, independent of any option quote.

- *Desk policy.* Very often a is simply fixed—to a value the desk trusts and holds across recalibrations—precisely so the daily fit reduces to the stable, exact $\sigma(t)$ bootstrap and the marks do not jump because the optimiser found a different local trade-off between a and $\sigma(t)$.

This a–$\sigma(t)$ identifiability trade-off is the central practical fact about calibrating the model: hold a down and the rest is easy and exact; let it float against a thin target and the fit turns soft and unstable.

The objective, when it is a least-squares fit. Where a (or a small handful of parameters) *is* fitted rather than fixed, the problem is a low-dimensional weighted least squares,

$$\min_{a,\,\sigma(\cdot)} \sum_k w_k \left(V_k^{\text{model}}(a, \sigma) - V_k^{\text{market}} \right)^2, \tag{21}$$

the sum running over the chosen quotes k. One fits to prices or to implied volatilities directly; the weights w_k encode *vega and liquidity*—a quote the book has real risk to, or that trades on a tight market, counts for more, and an illiquid corner of the grid counts for little or is dropped. The handful of parameters is then solved by Levenberg–Marquardt or any robust local optimiser. What makes it pleasant is the structure used throughout: each V_k^{model} is a closed form of §7, so the objective and its gradient are cheap and exact, the surface is smooth, and no simulation error can be mistaken for a poor fit [6]. In the common production setup—global a fixed by policy, $\sigma(t)$ bootstrapped—this outer optimisation collapses entirely into the exact sweep, and a full recalibration costs milliseconds.

Staging and discipline. A clean calibration runs in fixed stages, hardest constraint first: *(i)* fit the curve through $\alpha(t)$—exact, no choice involved; *(ii)* fix or fit a, the slow correlation parameter, and change it rarely; *(iii)* bootstrap $\sigma(t)$ to the chosen at-the-money strip. Two disciplines matter. First, keep $\sigma(t)$ *smooth*: a bootstrap that matches every quote to the basis point can produce a jagged, oscillating volatility that reprices today's strip perfectly and forecasts tomorrow's hedge ratios badly—a genuine overfitting danger. Lightly penalising the curvature of $\sigma(t)$, or fitting a smooth parametric form instead of free pieces, trades a little fit error for a lot of stability; with one factor it is the right trade. Second, match the recalibration cadence to the parameter: $\sigma(t)$ is refitted as often as the vol market moves—daily, or intraday on an active desk—while a is held across many recalibrations and revisited only deliberately, because a wandering a moves every mark and every hedge for no real informational gain. Curve exactly, a slowly, $\sigma(t)$ daily: that staging is what keeps a one-factor book's marks calm.

9 The smile the model makes, and the skew it cannot

Everything so far has been priced *at the money*: the drift matches today's curve, and a, σ are tuned to at-the-money cap and swaption volatilities. But options trade at many strikes, and their implied volatilities are not flat across strike—the market shows a *smile*, usually with a pronounced *skew*. What smile does the Hull–White model itself produce, and can it be bent to match the market's? The answer exposes a real limitation, and naming it is the rates analogue of the question that motivates smile modelling on the equity side.

Gaussian rates make a fixed smile. Because the short rate is Gaussian, so (to a close approximation) are forward rates; and a model with normal forward rates produces an almost *flat* smile once implied volatilities are quoted in *normal*, or Bachelier, terms (Figure 24). That is internally consistent and even elegant—but it is a *fixed* shape, dictated by the Gaussian assumption, not a free choice. The model has exactly two volatility knobs, a and σ, and they have already been spent on the level and term structure of at-the-money volatility (Section 8 above). Nothing is left to bend the smile with. A market that shows a real skew—receiver volatilities richer than payer volatilities, say—or extra curvature in the wings cannot be matched: the model has no parameter that moves implied volatility differentially across strike.

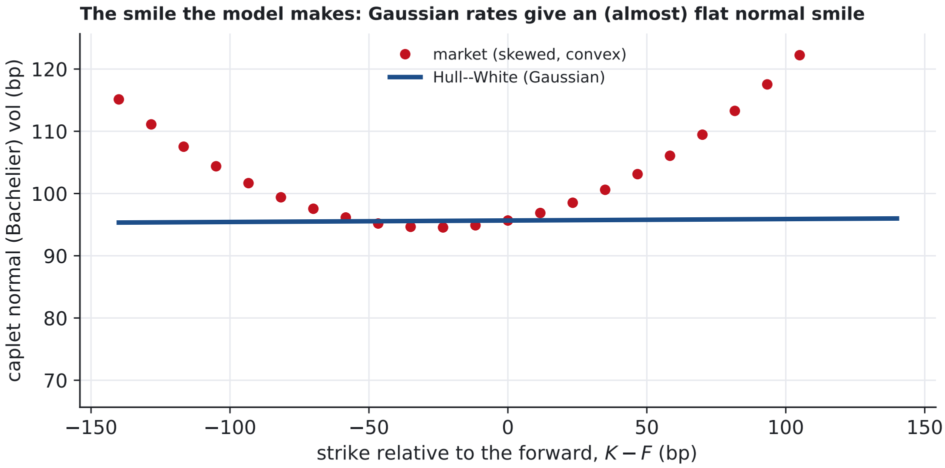

Figure 24: The smile the Gaussian model makes. Quoted in normal (Bachelier) vol, the one-factor caplet smile is almost flat across strike (blue): a fixed shape set by Gaussianity, with no free parameter to reshape it. A market smile (red) is skewed and convex; the plain model can match its level—the at-the-money point—but not its shape. Caplet at expiry 2y on the note's curve.

Two things "calibration" can blur. It helps to separate them. Hull–White fits the *at-the-money term structure*—volatility as a function of expiry—and does so well. It does *not* fit the *smile across strikes*—volatility as a function of moneyness at a fixed expiry—because Gaussianity fixes that shape. For products that live at or near the money, this is no obstacle. For products sensitive to the wings or the skew—deep out-of-the-money caps and floors, constant-maturity-swap convexity, anything with optionality far from the forward—it is a genuine gap.

How the market closes the gap. The standard extensions each add exactly the strike-dependence the plain model lacks, while keeping its spine—a short rate whose drift is slaved to the curve:

- *shifted (displaced) diffusion*—let the volatility scale with $(r + s)$ for a constant shift s, interpolating between normal ($s \to \infty$) and lognormal ($s = 0$) dynamics. This adds a controllable *skew*; it is the cheapest fix and a production standard.
- *local volatility on the short rate*—make the diffusion coefficient a function of the state, $\sigma \to \sigma(t, r)$, as in the quadratic-Gaussian and Cheyette families. This buys a full, calibratable smile while keeping much of the Gaussian tractability.
- *stochastic volatility*—give σ its own random process, as in SABR-type and stochastic-volatility Cheyette models. This generates a smile with realistic dynamics, at the cost of the closed forms.

All three keep the two-layer logic intact: the drift still pins the curve exactly, and the now-richer volatility parameters are still calibrated to option prices—there are simply more of them, and they now reach across strike as well as expiry. The plain Gaussian model is best read as the tractable core these smile-aware models extend, much as Black–Scholes is the core that local- and stochastic-volatility models extend on the equity side. That closes the account of what the model prices at and near the money—we turn next to the products that need a lattice rather than a formula.

10 Beyond closed form: Bermudans and callables

The closed forms of §7 cover the instruments one *calibrates* to. They do not cover the instruments one actually buys a term-structure model to *price*. Those are the products with early exercise or path dependence, and none of them has a formula.

Why a model is bought at all. A flat European swaption grid can be read straight off the market; you do not need a dynamic model to price what the market already quotes. You need one the moment the payoff depends on decisions taken *at several future dates*, or on the *path* the curve takes between now and then. A *Bermudan swaption* can be exercised on any of a schedule of dates, each time weighing immediate exercise against the right to wait. A *callable* or *puttable bond* embeds the same right to redeem early. A *range accrual* pays only on days the rate sits inside a band, accumulating along the path; a *target-redemption note* keeps paying coupons until their running total hits a cap, then unwinds. In every case the value depends on the joint distribution of future rates, not on a single terminal marginal—which is exactly what a calibrated dynamic model provides, and a quoted price cannot.

The lattice, and the early-exercise decision. The classic vehicle is Hull and White's trinomial tree [4] (Appendix B): a recombining lattice for the zero-mean factor x_t in which each node branches to three, with the branching *tilted*—down high up the tree, up low down—so the grid reproduces the mean reversion while staying recombining and efficient. The deterministic shift $\alpha(t)$ of (10) is layered on afterwards, so the finished tree reprices today's curve by construction. Pricing is then *backward induction*. Start at the final date, where every payoff is known, and roll back one step at a time. At each node the continuation value is the discounted, probability-weighted average of the successor values—(33)—the worth of *not* exercising and holding the option one more step. For an early-exercise product, overwrite that node value with

$$V_{\text{node}} = \max\bigl(\text{exercise value},\ \text{discounted expected continuation}\bigr), \tag{22}$$

the larger of cashing in now and carrying on. Sweeping (22) from maturity back to today threads the optimal exercise policy through the whole tree, and the value at the root is the price. This single comparison, node by node, is the heart of pricing a Bermudan.

Monte Carlo when the state grows. The tree is unbeatable in one or two dimensions but its node count explodes as the state grows. Once the payoff carries its own memory—path-dependent coupons, an accrued range, a redemption target—or once several factors drive the curve, one simulates the decomposition (10) forward instead and prices by Monte Carlo. The early-exercise decision, which a forward simulation cannot make directly because the continuation value lies in the future, is recovered by the Longstaff–Schwartz method: at each exercise date, regress realised continuation values on a few functions of the current state to *estimate* the continuation value, then apply the same rule (22) pathwise. It is backward induction's decision, fitted by regression rather than read off a grid.

The products, with numbers. It helps to see the machinery actually used. The four families below are the staple book of a rates-exotics desk. Every number that follows is computed on *this note's* curve and Hull–White dynamics ($a = 0.05$, $\sigma = 1\%$): the analytic pieces by the closed forms of §7, the early-exercise pieces by the trinomial tree, and the path-dependent ones by Monte Carlo—each cross-checked against an independent calculation. Notionals are EUR 100m; values are quoted in basis points of notional, or in price points (par = 100) for the notes.

1. Bermudan swaption. The flagship, and the reason the tree exists. A Bermudan payer swaption is the right to enter *one* fixed-maturity payer swap on any of a schedule of dates, choosing the moment. It is the option hiding inside almost every other product here—callable bonds and cancellable swaps embed exactly it—so a desk calibrates it to the co-terminal European swaptions of §8 and prices it by backward induction. Its value exceeds that of every constituent European, because the holder is not locked into a single expiry; the excess is the *switch* (or early-exercise) premium, and it is precisely what no market quote contains.

Bermudan payer swaption — term sheet

Notional	EUR 100m
Type	payer: pay fixed K, receive floating
Fixed rate K	3.21% (the 10y par swap rate, so at-the-money)
Final maturity	10 years
Exercise	annually, years 1–9; on exercising at year e, enter the payer swap running from e to year 10

Valuation. The nine co-terminal "$e \times (10-e)$" European swaptions, priced by Jamshidian (Appendix E), run from 274 bp at 1×9 up to 323 bp at 3×7 and down to 67 bp at 9×1 (Figure 25). Rolling the exercise rule (22) back through the tree—at every node and every exercise year, comparing immediate exercise against the discounted continuation—gives a **Bermudan value of 459 bp**, or EUR 4.59m. That is 136 bp *above* the most expensive European: the price of choosing *when* to exercise, which the model supplies and the market cannot.

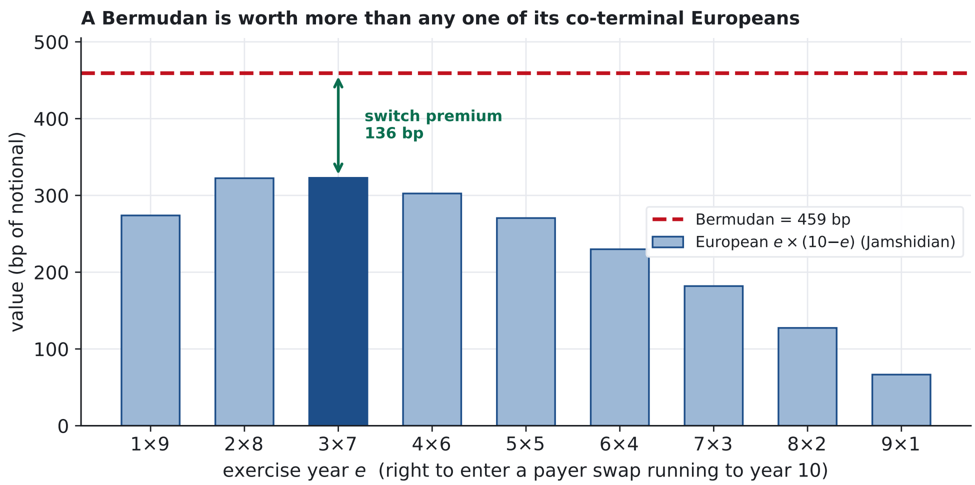

Figure 25: The flagship worked example. Each bar is one co-terminal European payer swaption (the Jamshidian closed form); the Bermudan, priced by backward induction on the tree, sits above all of them. The gap to the best single European—here 136 bp over the 3×7—is the switch premium, the value of the right to pick the exercise date, and the one number on the page that no European quote contains.

2. Callable bond and cancellable swap. The same option, sold by an issuer. A *callable bond* lets the issuer redeem early, at par, on a call schedule (a "10 nc 3" is a ten-year bond, non-call for three years); the investor is *short* a Bermudan and is paid a fatter coupon for it. A *cancellable swap* embeds the mirror right—the payer may tear the swap up early—which is a Bermudan *receiver* swaption. Both split into a vanilla leg minus that Bermudan, so the desk prices and hedges them with the engine above.

Callable bond ("10 nc 3") — term sheet	
Notional	EUR 100m
Maturity	10 years
Coupon	fixed, annual
Issuer call	annually at par (100), years 3–9
Redemption	par at call or at maturity

Valuation. Off the curve the straight 10y par yield is 3.206%. At a 4% coupon the *non-callable* bond is worth 106.70, but the issuer's embedded call—a Bermudan, rolled back on the tree—is worth 6.59 points, leaving the **callable bond at** 100.12. To sell it at exactly par the issuer must lift the coupon to **3.975**%: a 77 bp step-up over the straight par yield, which is the investor's pay for being short the call. *(The mirror trade: the right to cancel a 10y par payer swap from year* 3 *is a receiver Bermudan worth* 408 *bp, equivalent to a* 48 *bp step-up in the swap's breakeven fixed rate.)*

3. Range accrual note. A yield play, and the model's digital engine at work. The coupon accrues only on the days a reference rate sits inside a band, so the holder is *short a strip of digital options* on the rate and is paid a high headline coupon for the days it gives up. The valuation rests on one fact from §7: a Gaussian short rate makes the future rate a monotone function of one state, so each day's "in-the-band" digital has a closed form; the coupon is their discounted sum, and any call feature is layered on by the same backward induction.

Range accrual note — term sheet	
Notional	EUR 100m
Maturity	5 years
Coupon	5% × (days the 3m rate fixes in [0%, 4%])/(days in period), paid annually
Redemption	par at maturity

Valuation. Summing the daily digitals, the expected accrued fraction falls year by year—92%, 74%, 59%, 65%, 62%—as the forward 3m rate drifts towards the 4% cap and its distribution widens. The coupon stream is worth 16.16 points, so with principal the **note is worth** 101.34, against 107.98 for an unconditional 5% note: a 6.6-point (664 bp) give-up, which is the value of the digitals the investor has sold.

4. CMS and target-redemption notes. Two more reasons for a model: convexity, and path. A *constant-maturity-swap* payoff references a long swap rate paid on the "wrong" schedule; its expected fixing therefore carries a *convexity adjustment* that only a model can size. A *target-redemption note* (TARN) pays a coupon—often inverse-floating—and redeems early the moment the cumulative coupon hits a cap: pure path dependence, so it is priced by Monte Carlo with the early-redemption test applied along each path.

CMS caplet (on the 5y swap rate, fixing at year 5, paid then), *and* a **TARN** (10y, annual coupon max(8% − 1y rate, 0), target total 12%).

CMS valuation. The forward 5y swap rate is 3.153%; the model's convexity adjustment lifts the expected fixing to 3.250% (+9.7 bp), and an at-the-money CMS caplet is worth 66 bp.

TARN valuation. Monte Carlo over 80,000 paths (the inverse-floating coupon is high when rates are low, so the target is hit faster in low-rate paths): **fair value** 101.6 **points**, expected life 3.2 years, with a 99% chance of redeeming by year 5. The standard error is under 0.02 of a point.

Where one factor starts to bite. Notice what the exercise rule depends on. At each date the decision compares the value of exercising—which is the value of the remaining swap or bond, a function of the *whole curve* seen from that date—to the value of waiting. A Bermudan is therefore sensitive not just to where rates are but to how rates of different maturities *co-move*: the shape of the curve at each decision date, and how that shape can twist between dates. A single Gaussian factor moves the entire curve in lockstep, leaving it only one shape to twist into—and that rigidity is precisely the limitation the rest of this part turns to (§11, and the two-factor remedy of Part II).

11 What one factor cannot do

The one-factor model has a single source of randomness, $\mathrm{d}W_t$. Every rate on the curve is driven by that one shock, scaled by its own loading $\sigma e^{-a(T-t)}$—the single fixed profile we met back in Figure 16. It is worth being precise about what this does and does not buy. The decaying loading gives a genuine *volatility term structure*—the short end moves more than the long end (Figure 17)—so the loadings differ in *size*. But a difference in size is not a difference in *direction*: the move is still the one shock, so when the short rate goes up every other rate goes up too, just by a smaller amount. Two forward rates, however far apart in maturity, are therefore *perfectly correlated* in a one-factor model: they never move in opposite directions, only in lockstep. The curve can only move up and down and pivot rigidly about the mean reversion; it cannot genuinely *twist*, with the front end falling while the long end rises. Figure 27 (the flat line at correlation 1) makes the point that Part II's second factor is built to fix.

For many products this does not matter: a Bermudan swaption into a single fixed swap, a callable bond, a cap—all depend essentially on the level of one rate or a tightly clustered set, and a well-calibrated one-factor model prices them well and cheaply. It is the first choice precisely because it is so tractable. The limitation bites for products whose value *is* the relative movement of two parts of the curve: options on the slope (CMS spread options), some long-dated Bermudans whose exercise boundary depends on the shape of the curve, and any structure marked against the *decorrelation* of rates. For those, one factor is not enough—not because it fits today's curve any less exactly, but because its *dynamics* are too rigid. That is the cue for a second factor.

Part II
The two-factor model

12 Why a second factor

Adding a second random driver does three related things, all of them about dynamics rather than the static fit (which remains exact in the same way).

It decorrelates the curve. With two independent-ish shocks, a near rate and a far rate need no longer move together: each is now a different blend of the two drivers, so their correlation falls below one and decays as the maturities separate. The simplest way to picture the change is to plot one rate's move against another's (Figure 26). In a one-factor model the two are the same shock rescaled, so the scatter collapses onto a line—correlation exactly one. A second, independent shock fattens that line into a cloud, and that cloud *is* a correlation below one: the front end can now fall while the long end rises, which the line could never show. This is the headline cure. Empirically the correlation between, say, the 2-year and 10-year rates is high but distinctly below one; Figure 27 quantifies how it decays as the maturities separate, and a realistic model has to reproduce that to price anything sensitive to the slope.

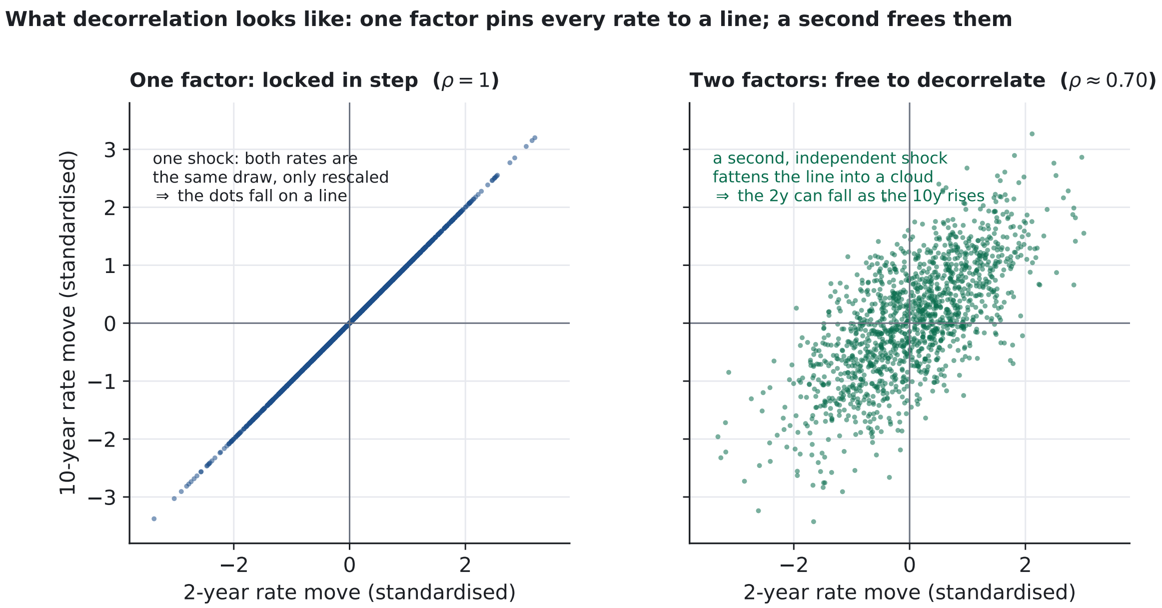

Figure 26: What decorrelation *is*, in one picture. Each dot is one scenario: the move in the 2-year rate (horizontal) against the move in the 10-year rate (vertical), standardised so that only the *shape* of the co-movement shows, not its size. *Left:* one factor—a single shock drives both, so every rate is that one draw rescaled and the dots collapse onto a line (correlation exactly 1). *Right:* two factors—a second, independent shock lets the two maturities move by different amounts *and* in different directions, fattening the line into a cloud (correlation about 0.7). The volatility term structure sets how *big* each rate's move is; decorrelation is a separate property, and it is the one that needs the second factor.

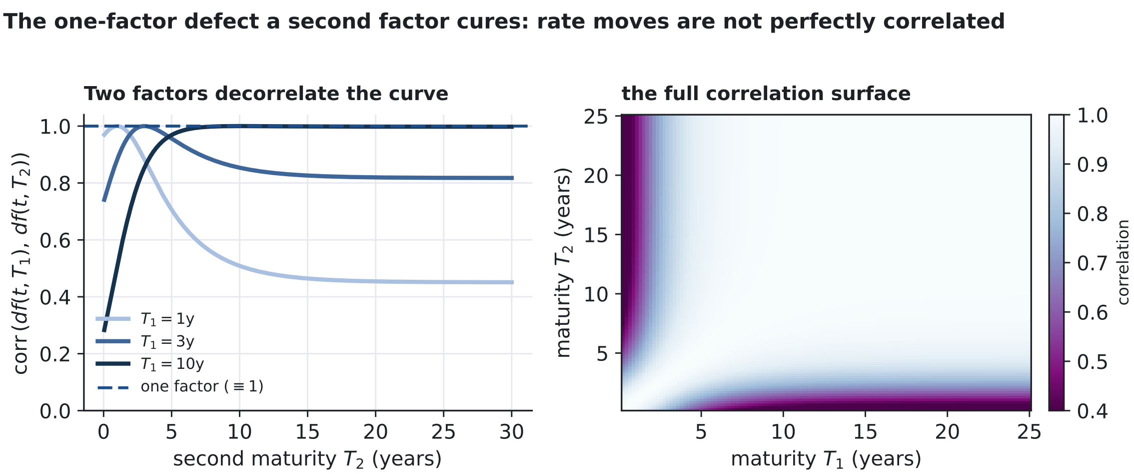

Figure 27: The defect, and the cure. *Left:* the instantaneous correlation between forward-rate moves at maturities T_1 and T_2. In the one-factor model it is identically 1 (dashed)—all rates move together. The two-factor model (solid) lets it fall away as the maturities separate, the more so for shorter T_1. *Right:* the full two-factor correlation surface, high near the diagonal (neighbouring maturities) and lower in the corners (front end versus long end).

It lets the curve twist. The two factors can be tuned to one fast and one slow mean reversion. The fast factor decays quickly along the curve, so it moves the *front end*; the slow factor reaches far out, so it moves the *whole level* (Figure 28, left). Their combinations reproduce the moves a curve actually makes—parallel shifts, steepenings, flattenings, twists (Figure 28, right)—rather than the single rigid mode of one factor. These are essentially the first two principal components, level and slope, of historical curve movements.

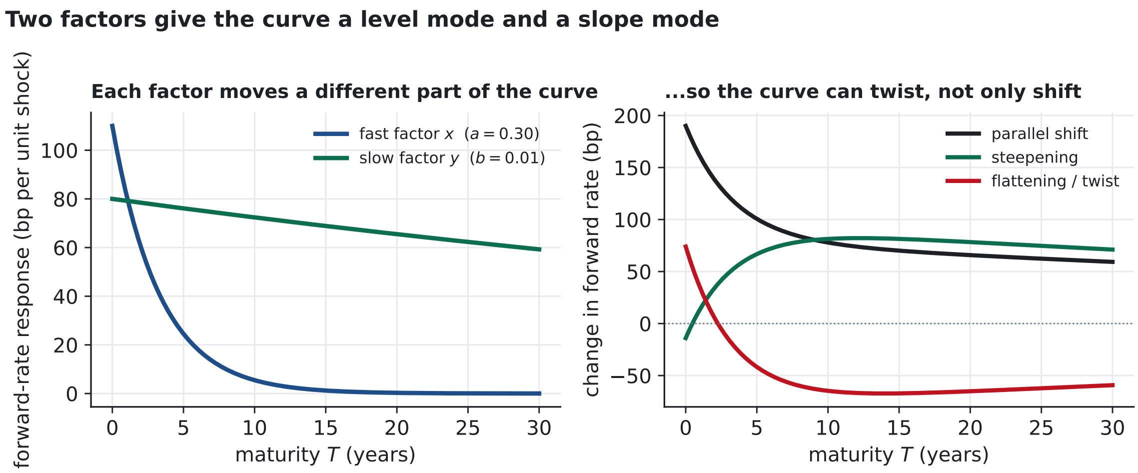

Figure 28: Two factors, two modes. *Left:* the response of the forward curve to each factor—a fast-reverting factor that moves mainly the front end (blue) and a slow one that moves the whole curve (green). *Right:* combining them generates shifts, steepenings and twists, the dominant modes of real curve movements, none of which a single factor can produce.

It fits the humped volatility term structure. Market at-the-money volatilities, plotted against option expiry, are typically *humped*—rising to a peak at a year or two, then falling. A single constant (a, σ) can only produce a monotone term structure (Figure 20), so it cannot fit a hump. Two factors with opposite-signed correlation can: the cross term suppresses volatility at the very short end and the surviving slow factor lifts it further out, bending the term structure into a hump (Figure 29).

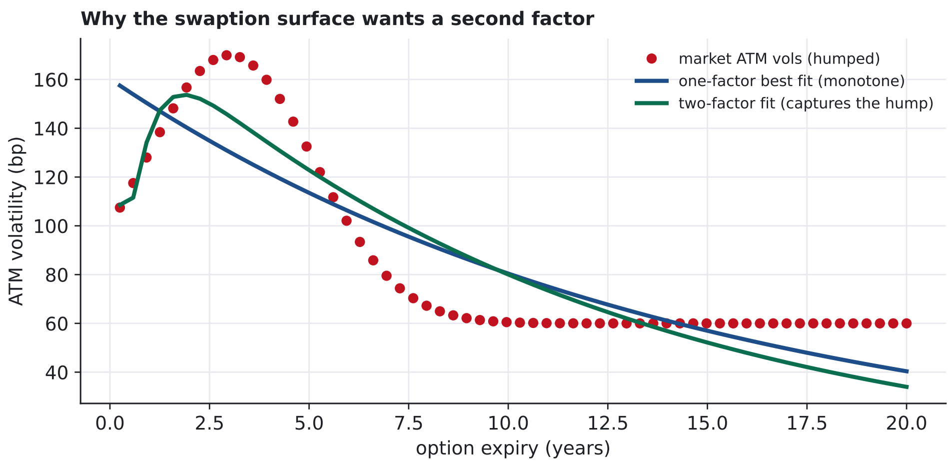

Figure 29: Fitting a humped market volatility term structure. A one-factor model, constrained to a monotone shape, can only split the difference (blue). The two-factor model reproduces the hump (green), because two timescales and a negative factor correlation give the term structure genuine curvature. (Illustrative least-squares fits to a synthetic humped market.)

13 The two-factor (G2++) model

The cleanest formulation of the two-factor Gaussian model is the additive *G2++* form [6], the direct analogue of the $r = x + \alpha$ decomposition (10). The short rate is the sum of two zero-mean

Ornstein–Uhlenbeck factors and a deterministic shift,

$$r_t = x_t + y_t + \varphi(t), \qquad \begin{aligned} dx_t &= -a\, x_t\, dt + \sigma\, dW_t^1, \\ dy_t &= -b\, y_t\, dt + \eta\, dW_t^2, \end{aligned} \qquad dW_t^1\, dW_t^2 = \rho\, dt. \tag{23}$$

The two factors have their own mean-reversion speeds a, b and volatilities σ, η, and—crucially—their Brownian shocks are correlated with coefficient ρ. (This is the same model as two-factor Hull–White; G2++ is just the tidiest parameterisation, and is equivalent to it up to a change of variables.)

The shift φ, and where it comes from. As in one factor, $\varphi(t)$ is not a free function: it is pinned by the single demand that the model reprice today's curve, $P(0,T) = \mathbb{E}\big[e^{-\int_0^T r}\big]$ for every T. The argument is the one that fixed α in §5 (eq. (12)), only now the random part of the short rate is the *sum* $x_t + y_t$ of two correlated Gaussians. Matching the curve asks again for "forward rate plus a Jensen convexity bump," and the bump is the variance of that combined driver. Because the variance of a sum is $\operatorname{Var} x + \operatorname{Var} y + 2\operatorname{Cov}(x,y)$, it splits into each factor's *own* convexity plus a cross term—the covariance, present only because the shocks are correlated:

$$\varphi(t) = f(0,t) + \underbrace{\frac{\sigma^2}{2a^2}\big(1-e^{-at}\big)^2}_{\text{factor 1's own bump}} + \underbrace{\frac{\eta^2}{2b^2}\big(1-e^{-bt}\big)^2}_{\text{factor 2's own bump}} + \underbrace{\rho\,\frac{\sigma\eta}{ab}\big(1-e^{-at}\big)\big(1-e^{-bt}\big)}_{\text{covariance / cross term}}. \tag{24}$$

Term by term: $f(0,t)$ is the forward rate, carrying the curve's shape exactly as before; the first bump, $\frac{\sigma^2}{2a^2}(1-e^{-at})^2$, is *identical in form* to the one-factor convexity correction of (12), and the second is its twin for the slower factor; and the last term is the covariance of the two, which with $\rho < 0$ is negative and so *trims* the combined bump (the factor of two in $2\operatorname{Cov}$ having cancelled the $\frac{1}{2}$ that accompanies a Jensen correction). The exact-fit property is untouched: whatever the five dynamics parameters $(a, b, \sigma, \eta, \rho)$ are, $\varphi(t)$ recomputes itself to reprice the curve. Everything new lives in the dynamics those five parameters control.

Still Gaussian, still affine. The sum of two Gaussian factors is Gaussian, so the model keeps the conveniences of Part I. Bond prices remain affine, now in two state variables,

$$P(t,T) = A(t,T)\, \exp\big(-B_a(t,T)\, x_t - B_b(t,T)\, y_t\big), \qquad B_a(t,T) = \frac{1-e^{-a(T-t)}}{a}, \tag{25}$$

with B_b defined analogously with b. The coefficient $A(t,T)$ is the exact analogue of the one-factor A of (15)—today's curve, dressed by the model's accumulated variance:

$$A(t,T) = \frac{P(0,T)}{P(0,t)}\, \exp\Big(\tfrac{1}{2}\big[\, V(t,T) - V(0,T) + V(0,t)\,\big]\Big), \tag{26}$$

where $V(t,T) = \operatorname{Var}\big[\int_t^T r_u\, du\big]$ is the variance of the integrated short rate—built from the same two factor volatilities and their correlation, so it carries the same three-part structure (writing $\tau = T - t$):

$$\begin{aligned} V(t,T) = {} & \frac{\sigma^2}{a^2}\Big[\tau + \tfrac{2}{a}e^{-a\tau} - \tfrac{1}{2a}e^{-2a\tau} - \tfrac{3}{2a}\Big] + \frac{\eta^2}{b^2}\Big[\tau + \tfrac{2}{b}e^{-b\tau} - \tfrac{1}{2b}e^{-2b\tau} - \tfrac{3}{2b}\Big] \\ & + \frac{2\rho\sigma\eta}{ab}\Big[\tau + \tfrac{e^{-a\tau}-1}{a} + \tfrac{e^{-b\tau}-1}{b} - \tfrac{e^{-(a+b)\tau}-1}{a+b}\Big], \end{aligned} \tag{27}$$

the first two brackets each factor's own accumulated variance and the third their covariance—fast, slow and interference once more. It is long, but every piece is a closed form in the five parameters; and like its one-factor parent it still reprices the curve whatever they are—set $t = 0$ and the bracket collapses ($V(t,T) - V(0,T) \to 0$ and $V(0,t) \to 0$), leaving $A(0,T) = P(0,T)$.

The forward-rate volatility, and its hump. The volatility term structure comes from differentiating the bond, exactly as it did in one factor (§6, eq. (16)). Since $f(t,T) = -\partial_T \ln P(t,T)$ and $\ln P = \ln A - B_a(t,T)\,x_t - B_b(t,T)\,y_t$,

$$f(t,T) = -\partial_T \ln A + \big(\partial_T B_a\big)\,x_t + \big(\partial_T B_b\big)\,y_t,$$

and because $\partial_T B_a(t,T) = e^{-a(T-t)}$ (and likewise $e^{-b(T-t)}$ for B_b) the random part of the T-forward rate is the loaded sum $e^{-a(T-t)}\,x_t + e^{-b(T-t)}\,y_t$. Its instantaneous variance is, once more, a variance-of-a-sum—two own terms and a cross term (writing T for the time to maturity):

$$\nu(T)^2 = \underbrace{\sigma^2 e^{-2aT}}_{\text{fast factor}} + \underbrace{\eta^2 e^{-2bT}}_{\text{slow factor}} + \underbrace{2\rho\,\sigma\eta\,e^{-(a+b)T}}_{\text{interference}}. \tag{28}$$

Term by term: $\sigma^2 e^{-2aT}$ is the fast factor's contribution, its shock σ damped by e^{-aT}, so it moves mostly the *short* end; $\eta^2 e^{-2bT}$ is the slow factor's, decaying gently and reaching *far* out the curve; and $2\rho\sigma\eta\,e^{-(a+b)T}$ is the interference between the two. That cross term is the crux: with $\rho < 0$ it *subtracts*, and most strongly where both factors are still alive—the short-to-mid maturities—scooping volatility out there and leaving a *hump* further out (Figure 29). The very same blend, taken as a *ratio* between two maturities, is the instantaneous correlation of Figure 27. These two properties—a volatility term structure flexible enough to hump, and a maturity-dependent correlation below one—are the entire reason the model is worth its five parameters.

14 Pricing and calibration with two factors

What survives, what changes. Caps and floors are still essentially options on a single rate, and they keep closed-form (Gaussian bond-option) prices under (25). European swaptions are the casualty of the second factor: Jamshidian's trick relied on every bond price being monotone in *one* state, and with two states the exercise boundary is a *curve* in the (x, y) plane rather than a single point. The standard resolution keeps the model practical: condition on one factor, price the resulting one-dimensional problem analytically, and integrate over the other factor numerically—a single well-behaved one-dimensional quadrature, fast enough to sit inside a calibration loop. European swaptions are thus *semi-analytic*: not a clean formula, but not a simulation either. Bermudans, callables and path-dependent structures go on a two-dimensional lattice or, more commonly, Monte Carlo on the two factors of (23).

Figure 30 is the picture of exactly what changed. In one factor the swaption's exercise decision was a single threshold on the rate line (Figure 21); with two factors the underlying swap depends on both x and y, so the set where exercise is worthwhile is a half-plane bounded by a *curve* in the (x, y) plane, not a point. There is no single crossing rate for Jamshidian to hang the decomposition on—hence the fall back to a one-dimensional numerical integration over the cloud.

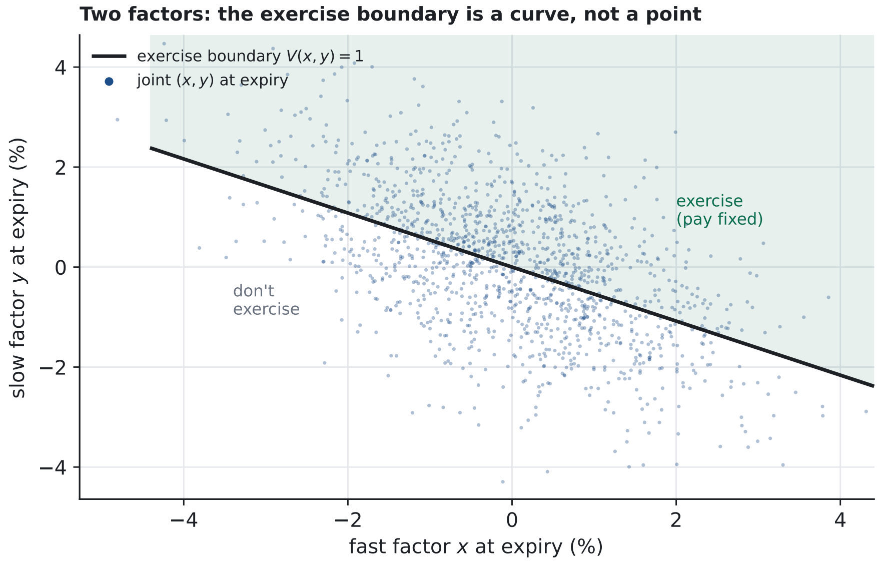

Figure 30: Why two factors break the closed-form swaption. The joint distribution of the two factors at expiry (blue cloud, tilted by their negative correlation) sits in a plane, and the exercise region is bounded by a *curve* $V(x, y) = 1$, not a single threshold. Jamshidian's decomposition needed the one crossing point of Figure 21, which no longer exists; the price falls back to a one-dimensional integration over this plane.

Calibrating five parameters: a genuine fit, not a bootstrap. This is where the second factor exacts its price. Recall how clean one-factor calibration was (§8): with a fixed, the single $\sigma(t)$ *bootstraps* exactly along a strip—one quote pins one σ_i, swept outward, no optimiser and no residual. None of that survives. The five dynamics parameters $(a, b, \sigma, \eta, \rho)$ enter the swaption price *together*, through the vol blend (28) and its correlation; there is no triangular structure to unwind, so calibration becomes a genuine multi-dimensional least-squares fit. Three things change in kind, not merely in degree.

The target must be two-dimensional. A one-factor fit lives on a single strip of volatilities, one per expiry; but a strip says nothing about how rates of *different* maturities co-move, which is the very property the second factor exists to set. So a two-factor fit needs the whole swaption *grid*—expiries down one side, swap tenors across the top (Figure 22)—because it is the variation *across tenor*, at a fixed expiry, that exposes the curve's internal correlation. Fit only a strip and ρ is invisible: the model will match that strip exactly and still misprice everything slope-sensitive.

The roles separate, but two of them entangle. The five parameters keep near-distinct jobs, which is what makes the larger fit tractable at all:

- σ, η set the overall volatility *levels* of the two factors;
- a, b—one fast, one slow—set the two *timescales*, hence the shape of the volatility term structure (the hump of Figure 29) and how far along the curve each factor reaches;
- ρ sets the *correlation* between curve points, hence the decorrelation of Figure 27 and the pricing of every slope-sensitive instrument.

The awkwardness is that the single parameter ρ does *two* jobs at once: a negative ρ both bends the vol term structure into its hump *and* decorrelates the curve. The hump and the correlation are thus fit jointly, and only data that pins each separately—a humped vol term structure *and* a correlation-sensitive quote (the cross-tenor grid, or a CMS-spread mark)—can tell ρ's two effects apart. With a and b also free to trade off against the volatilities, the objective has shallow valleys, and an optimiser handed all five at once will happily wander into a different trade-off each day.

So the fit needs discipline. The recipe is to stage it along those near-orthogonal roles rather than fit five parameters jointly: fix the mean reversions a, b to stable, structural values (one fast, one slow, held across recalibrations); fit σ, η and a mild time-dependence to the vol levels and hump along the expiry axis; and let the cross-tenor grid or a CMS-spread mark pin ρ. Resist over-fitting the illiquid corners of the grid, where a fifth parameter is only too glad to chase noise. The one mercy is speed: the swaption price stays semi-analytic—a single quadrature—so even this larger, coupled fit runs noise-free and fast enough for a daily mark, exactly as in one factor.

15 What two factors price that one cannot

The decorrelation, twist and hump of §12 are not abstract virtues; they are what let the model price a whole class of products that one factor gets *qualitatively* wrong. The common thread is the *spread*: any payoff that depends on the gap between two points on the curve. One factor moves every rate off a single shock, so all rates are perfectly correlated and a spread has almost no volatility of its own; the second factor gives the spread a life. The four examples below are priced—like those of §10—on this note's curve, now with the G2++ dynamics of §13, by Monte Carlo over the two correlated factors. To make the failure unmistakable, each is priced *twice*: in the two-factor model, and in a one-factor model *calibrated to the very same single-rate vols*—matched to both the 2y and the 10y swaption, so the two models agree on every vanilla and differ only in correlation.

1. CMS spread option. The flagship, and the cleanest test. A *constant-maturity-swap spread* option pays on the difference of two swap rates—say the 10-year minus the 2-year rate—observed at one future date. Its entire value is the *volatility of that spread.* In one factor the 10y and 2y rates are the same shock rescaled, perfectly correlated, so the spread barely moves and the option is worth almost nothing; the second factor decorrelates them and gives the spread a real volatility (Figure 31).

CMS spread option — term sheet

Notional	EUR 100m
Underlying	10y swap rate − 2y swap rate, fixing in 5 years
Expiry	5 years
Payoff	$N\,(S_{10} - S_2 - K)^+$ at year 5, struck at the forward spread

Valuation. Both models are matched to the same 2y and 10y swaption vols (1.54% and 1.53%), and to the same forward spread (−0.03%); the only thing left to differ is the correlation, 0.93 in two factors versus 1.00 in one. That alone collapses the spread volatility from 56 bp to under 1 bp, and with it the option:

$$2\text{F} = 19.1 \text{ bp}, \qquad 1\text{F} = 0.3 \text{ bp} \quad (\text{a } 60\times \text{ gap}).$$

Every basis point of the difference is spread volatility the one-factor model cannot make.

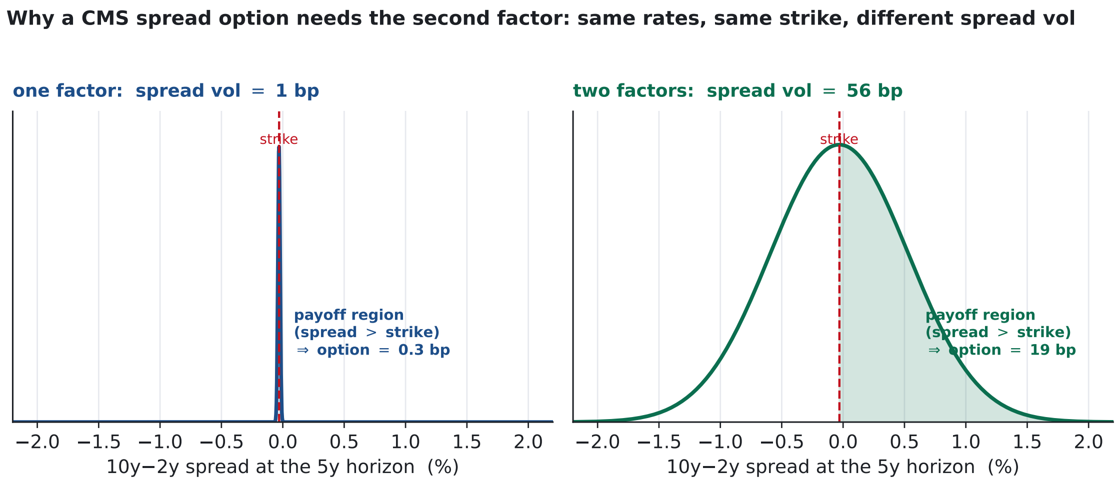

Figure 31: Why the spread option needs the second factor. The distribution of the $10y - 2y$ spread at the five-year horizon, in two models calibrated to the *same* 2y and 10y volatilities. *Left:* one factor forces the two rates to move in lockstep, so the spread is all but deterministic (vol 1 bp)—almost no mass clears the strike, and the option is worth 0.3 bp. *Right:* two factors decorrelate the rates, the spread acquires a real 56 bp volatility, and the shaded payoff region—hence the option, at 19 bp—comes to life. Same rates, same strike; the entire difference is the spread's volatility.

2. CMS spread range accrual. The digital cousin. The coupon accrues only on the days the $10y - 2y$ spread sits inside a band—a strip of digital options on the spread. Because one factor holds the spread essentially fixed, it is *certain* the spread stays put: it accrues a near-constant fraction and badly *over*-values the coupon. Two factors let the spread diffuse out of the band, giving the realistic, declining accrual.

CMS spread range accrual — term sheet

Notional	EUR 100m
Maturity	5 years
Coupon	$4\% \times$ (days the $10y - 2y$ spread is within $\pm 0.5\%$)/(days), annual
Redemption	par at maturity

Valuation. In two factors the spread wanders out of the band as time passes, so the expected accrual falls year by year—$89\%, 68\%, 63\%, 63\%, 63\%$—for a coupon stream worth 12.7 points. In one factor the spread is frozen inside the band, so it accrues 100% every year and the coupon is worth 18.2 points: an over-valuation of 5.5 points, the entire amount being spread movement the one-factor model has assumed away.

3. Steepener note. The starkest illustration. A steepener pays a leveraged coupon on the curve's slope, $\max\big(g\,(S_{10} - S_2),\, 0\big)$—a bet that the curve steepens. On this (flat-to-inverted) forward curve the one-factor spread is deterministic and slightly negative, so $\max(g \cdot \text{spread}, 0)$ is zero on *every* path and the coupon is worth *exactly nothing.* The two-factor spread sometimes turns positive, and the coupon comes alive.

Steepener note — term sheet

Notional	EUR 100m
Maturity	5 years
Coupon	$\max\big(10 \times (S_{10} - S_2),\ 0\big)$, paid annually
Redemption	par at maturity

Valuation. 2F coupon stream = 6.35 points; 1F coupon stream = 0.00. One factor does not merely under-price the steepener—it cannot price it at all, because a model in which the curve never changes shape sees a payoff on changing shape as worthless.

4. Bermudan swaption, re-examined. The counterpoint, and a guard against over-reach. The Part I Bermudan (§10) is *not* a spread product: its value is driven by the *level* of one swap rate at each exercise date, not by the gap between two. Re-priced in two factors—both models again calibrated to the same 2y and 10y vols—it is 368 bp, against 372 bp in one factor: a move of under 1%. (The level differs from the 459 bp of §10 because both models here are tied to the flatter two-factor volatilities, not that section's illustrative pair; it is the two-factor-versus-one *difference* that matters.) The second factor that *transforms* the spread products barely touches the Bermudan—which is exactly why one reaches for it by product, not by reflex.

The pattern is the model-choice table of §16 in miniature: a second factor earns its keep precisely when the payoff sees the *shape* of the curve—its slope, or the correlation between its points—and not a moment before.

16 Choosing a model, and what lies beyond

The choice between one and two factors is a question about *dynamics*, never about the static fit—both reprice today's curve exactly. It turns on what the product's value depends on (Table 1).

product	dominant sensitivity	natural model
caps, floors, European swaptions	level of one rate	one factor (analytic)
callable/puttable bonds, single-callable Bermudans	level of one rate, early exercise	one factor
long-dated/complex Bermudans, where the exercise region depends on curve shape	level + some slope	one or two factors
CMS spread options, curve steepeners, slope-sensitive structures	correlation / slope of the curve	two factors
products needing a rich vol term structure (humped)	shape of the vol surface	two factors

Table 1: A rough guide to model choice. The split tracks one question: does the value depend on the *level* of one rate (one factor suffices), or on the *relative movement* of different parts of the curve (a second factor is needed)?

The rule of thumb. Reach for one factor by default—it is analytic, fast, easy to calibrate, and right for the large class of products driven by a single rate. Reach for two factors when the payoff sees the *shape* of the curve: its slope, the correlation between its points, or a volatility term structure with a hump that one factor cannot bend to. A third factor (curvature) is occasionally added for the most curve-sensitive books, with rapidly diminishing returns and rising calibration fragility.

What lies beyond. The Hull–White family is the Gaussian, short-rate corner of a larger world. Its limits are the obvious ones: rates are Gaussian (a clean smile is hard, and there is no skew without extending the model), and the short rate is a slightly artificial object to model directly. The market-standard alternatives address each. *Market models* (the LIBOR/forward market model of Brace–Gatarek–Musiela [8]) model the observable forward rates directly and naturally accommodate the full smile across many factors, at a much higher computational cost. For the volatility smile specifically, one adds stochastic volatility or a local-volatility-style mapping to the short-rate factors (shifted-SABR dynamics, or a stochastic-volatility Cheyette model). But for a great deal of real-world interest-rate work—discount-curve-consistent pricing of callables, Bermudans and curve-sensitive exotics—one- and two-factor Hull–White remain the workhorses, for exactly the reason this note began with: they fit the whole curve exactly, and they stay simple enough to calibrate and compute with every day.

A From the short rate to the bond price

The affine formulas (14)–(15) are not guessed; they fall out of a pair of ordinary differential equations. Under the risk-neutral measure the bond price $P(t,T) = P(t,T,r_t)$ satisfies the pricing equation

$$\frac{\partial P}{\partial t} + \big(\theta(t) - ar\big)\frac{\partial P}{\partial r} + \tfrac{1}{2}\sigma^2\frac{\partial^2 P}{\partial r^2} - rP = 0, \qquad P(T,T) = 1, \tag{29}$$

the statement that the discounted bond price is a martingale. Try the *affine ansatz* $P(t,T) = \exp\big(A(t,T) - B(t,T)\,r\big)$, suggested by the linearity of the drift and the constant diffusion. Substituting and dividing by P gives

$$A_t - B_t\,r - \big(\theta(t) - ar\big)B + \tfrac{1}{2}\sigma^2 B^2 - r = 0. \tag{30}$$

This must hold for every r, so the terms linear in r and the terms constant in r vanish separately:

$$B_t = aB - 1, \qquad A_t = \theta(t)\,B - \tfrac{1}{2}\sigma^2 B^2, \tag{31}$$

both with terminal condition zero at $t = T$. The first is a linear ODE whose solution is $B(t,T) = (1 - e^{-a(T-t)})/a$—rising from 0 at maturity toward the reciprocal of the mean reversion. The second integrates A directly once B is known. Feeding in the curve-fitting $\theta(t)$ of (13) and simplifying produces the closed form (15) for $A(t,T)$ in terms of the initial discount curve and forward rate. (In the general affine framework B solves a Riccati equation; here, because the model is Gaussian, it is merely linear, which is the technical reason Hull–White is so tractable.) The two-factor bond price (25) comes from the same argument with two state variables: one linear B-equation per factor and a single A-equation that now also collects the cross term in ρ.

B The Hull–White trinomial tree

Early-exercise products are valued by backward induction on a lattice for the zero-mean factor x_t of (10); the shift $\alpha(t)$ is added to each node afterwards so the tree reprices the curve. Hull and White's [4] construction is a recombining *trinomial* tree: from each node the process branches to three, capturing the diffusion, while the *choice of which three* encodes the mean reversion (Figure 32).

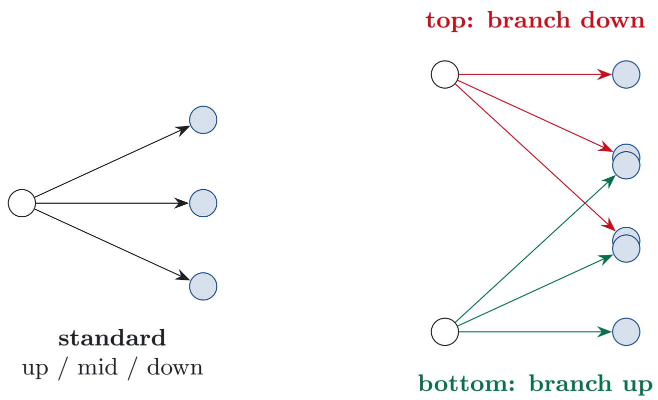

Figure 32: Branching in the Hull–White trinomial tree. In the bulk of the tree a node takes the *standard* up/mid/down branching (black). Near the top, where mean reversion would pull the rate back down hardest, the branching switches to down/mid/down-more (red); near the bottom it switches to up/mid/up-more (green). Choosing the branch pattern and probabilities this way reproduces the $-a\,x\,\mathrm{d}t$ drift on a recombining grid, keeping the tree both mean-reverting and efficient. Backward induction then prices early exercise: at each node, take the larger of exercising now and the discounted expected continuation value.

The probabilities on the three branches are chosen to match the mean and variance of the exact Gaussian increment of x_t over the step; switching the branch pattern at the top and bottom keeps those probabilities valid (between 0 and 1) even where a standard branching would imply a negative probability because the mean reversion is too strong. The result is a lattice that is exactly consistent with the model's first two moments, recombines (so its size grows linearly, not exponentially, with steps), and reprices the curve once α is layered on.

A worked step. Take $a = 0.10$, $\sigma = 1\%$ and a yearly step $\Delta t = 1$, so the grid spacing is $\Delta x = \sigma\sqrt{3\Delta t} = 1.73\%$. With the standard up/mid/down branching the probabilities at the node $x = j\,\Delta x$ are

$$p_u = \tfrac{1}{6} + \tfrac{a^2j^2\Delta t^2 - aj\Delta t}{2}, \quad p_m = \tfrac{2}{3} - a^2j^2\Delta t^2, \quad p_d = \tfrac{1}{6} + \tfrac{a^2j^2\Delta t^2 + aj\Delta t}{2}. \tag{32}$$

At the central node ($j = 0$) these are the symmetric $(p_u, p_m, p_d) = (\frac{1}{6}, \frac{2}{3}, \frac{1}{6}) = (0.167, 0.667, 0.167)$. One level up ($j = 1$) they tilt to $(0.122, 0.657, 0.222)$: the *down* branch is now the likeliest, which is precisely the mean reversion pulling the rate back toward the centre. (Far enough up, p_u in (32) would turn negative under standard branching—the level at which the tree switches to the downward-branching pattern of Figure 32.)

To price, attach the short rate $r_j = \alpha(t) + j\,\Delta x$ to each node and roll backwards. Suppose at the central node ($r = 3.0\%$) the three successor values one year out are $V_u = 98.0$, $V_m = 99.0$, $V_d = 100.0$. The node value is the discounted, probability-weighted average of its successors,

$$\begin{aligned} V &= e^{-r\Delta t}\big(p_uV_u + p_mV_m + p_dV_d\big) \\ &= e^{-0.03}\big(\tfrac{1}{6}\cdot 98.0 + \tfrac{2}{3}\cdot 99.0 + \tfrac{1}{6}\cdot 100.0\big) \\ &= e^{-0.03} \times 99.0 \\ &= 96.07. \end{aligned} \tag{33}$$

For an early-exercise product one simply overwrites V with $\max\big(V, \text{exercise value}\big)$ at each node before continuing back. That one line, (33), applied node by node from maturity to today, is the entire pricing algorithm.

C The forward measure and the Black formula

Two standard tools sit behind every closed-form price in the note: a change of measure that removes drift, and Black's formula that prices the resulting driftless option. We state both here so the main text can point back rather than repeat them.

The forward measure (change of numeraire). A price is an expectation—but of *what*, and under *which* measure? The general rule is that any positive traded asset may serve as the *numeraire*, the unit in which other prices are quoted: dividing every price by the numeraire turns those ratios into martingales under a measure attached to that choice. Taking the money-market account β_t as numeraire gives the risk-neutral measure of (5). For an option that settles on a single date T_O, the convenient choice is instead the T_O-bond, $P(\cdot, T_O)$—the T_O*-forward measure* $\mathbb{Q}^{T_O}$. Two things make it the natural home for a bond option.

First, it cuts the number of random quantities the price depends on from two to one. Under the risk-neutral measure the option, with payoff H at T_O, is worth

$$V_0 = \mathbb{E}^{\mathbb{Q}}\Big[\underbrace{\exp\big(-\textstyle\int_0^{T_O} r_s\,\mathrm{d}s\big)}_{\text{random discount factor}} \times \underbrace{H}_{\text{random payoff}}\Big], \tag{34}$$

an expectation of the *product* of two random and mutually correlated things: the stochastic discount factor, which depends on the whole rate path to T_O, and the payoff H, which depends on the underlying bond at T_O. Valuing it directly means integrating over the joint distribution of both. The change of numeraire dissolves the first of them. The general identity is that, under the measure $\mathbb{Q}^N$ attached to a numeraire N_t,

$$V_0 = N_0 \, \mathbb{E}^N\Big[\frac{H}{N_{T_O}}\Big], \tag{35}$$

and the T_O-bond is the inspired choice because it *matures at exactly the settlement date*, so $N_{T_O} = P(T_O, T_O) = 1$. Dividing by the numeraire at expiry then does nothing at all, while $N_0 = P(0, T_O)$ is merely today's bond price—a number known now. Equation (35) collapses to

$$V_0 = P(0, T_O) \, \mathbb{E}^{T_O}[\, H \,]. \tag{36}$$

The random discounting has been pulled clean out of the expectation and frozen into the single deterministic factor $P(0, T_O)$; what is left inside depends only on the payoff H, i.e. on the *one* underlying—the bond price (equivalently, the rate) at T_O. Two coupled sources of randomness have become one.

Second, that one remaining underlying is well behaved. The forward bond price $G(t) = P(t, T_B)/P(t, T_O)$—a ratio of a traded asset to the numeraire—is by construction a $\mathbb{Q}^{T_O}$-martingale: it has *no drift*, so only its variance Σ_P^2 enters the price. A driftless lognormal underlying and a clean discount factor are exactly what Black's formula consumes.

The formula. Black's 1976 formula prices a European option on a forward F that is lognormal at expiry, struck at K, settled and discounted by a factor D. The only volatility input is the *total* (terminal) volatility $v = \sigma\sqrt{T}$, the standard deviation of $\ln F$ accumulated to expiry. With

$$d_\pm = \frac{\ln(F/K) \pm \frac{1}{2}v^2}{v}, \tag{37}$$

the call and put are

$$\text{Call} = D\big[F\,N(d_+) - K\,N(d_-)\big], \qquad \text{Put} = D\big[K\,N(-d_-) - F\,N(-d_+)\big], \tag{38}$$

where N is the standard normal CDF. The two satisfy put–call parity, Call − Put $= D(F - K)$, as they must: a call minus a put is a forward. The shape is the familiar Black–Scholes one with the spot replaced by the forward and the discounting pulled out front—there is no drift term inside, because working with the forward F has already absorbed it.

Why a *Gaussian* model is priced with a *lognormal* formula. There is an apparent contradiction here worth clearing up, because it is the crux of why Black fits at all. Black's formula assumes the underlying F is *lognormal*; yet the Hull–White short rate is *Gaussian*—normally distributed, and even free to go negative. How can a normal-rate model be priced with a lognormal formula? The answer is that Black is applied not to the *rate* but to the *bond price*, and in an affine Gaussian model the bond price is the *exponential* of the rate. From the affine form (14),

$$P(t, T) = A(t, T)\, e^{-B(t,T)\, r_t}, \qquad \text{so} \qquad \ln P(t, T) = \ln A(t, T) - B(t, T)\, r_t.$$

The log bond price is an *affine function of the Gaussian* r_t, hence itself Gaussian; and a quantity whose logarithm is Gaussian is, by definition, *lognormal*. So the bond price—and with it the forward bond price $G = P(\cdot, T_B)/P(\cdot, T_O)$, whose log is also affine in r_t—is exactly lognormal (Figure 33). The lognormality Black requires is not an extra assumption layered onto the model; it is the direct consequence of the rate being Gaussian. *The exponential of a normal is a lognormal*—that is the whole story, and it is why the bond, not the rate, is the natural thing to hang the option on.

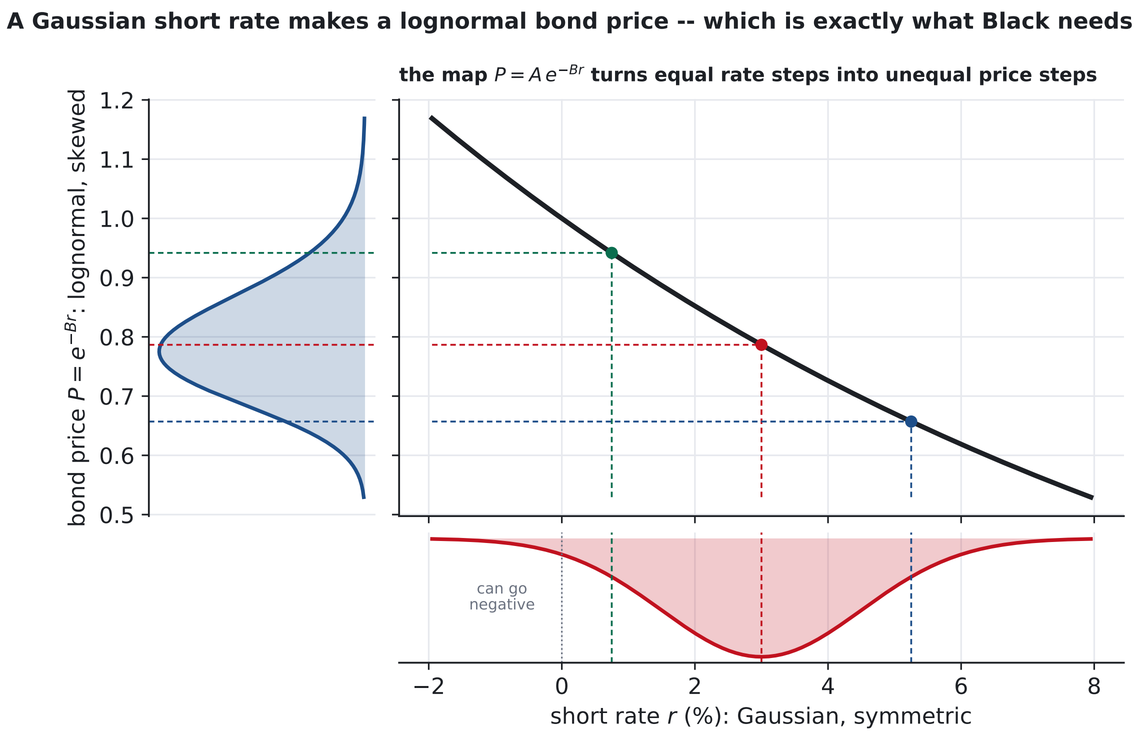

Figure 33: Why a Gaussian short rate is priced with a lognormal formula. The bond price is $P = A\,e^{-Br}$, the *exponential* of the (Gaussian) short rate. The symmetric bell of r (bottom, which may even stray below zero) is pushed through the decreasing, convex map (centre) into a *skewed, strictly positive* distribution of the bond price (left): a lognormal. Equal steps in the rate (the three dashed levels) become unequal steps in the price—the signature of the exponential. So the bond is lognormal precisely *because* the rate is normal, and Black's formula applies exactly.

There is a bonus in the word "exactly". Because the forward bond price's volatility Σ_P is *deterministic*—it depends on a, σ and the dates, not on the realised path—the forward bond price is lognormal with a *known* total variance, so Black's formula is an exact valuation of a zero-coupon bond option in this model, with both the lognormality and the volatility derived from the model rather than assumed of the underlying.

A last caution, to square this with market quotes. The *bond* is lognormal, but the forward *rate* $F = \frac{1}{\tau}\big(1/P(T_O, T_B) - 1\big)$ is not: it is an affine function of a lognormal, and sits closer to *normal* than to lognormal. That is exactly why the natural quoting convention for these options is the *normal* (Bachelier) volatility rather than Black's lognormal one, and why the model's caplet smile is flat in normal terms (§9). Black prices the *bond* option exactly; whether one then reports the answer as a lognormal-rate vol or a normal-rate vol is a convention applied after the fact.

Why this is the whole toolkit. Each analytic price in the note is an instance of (38), got by reading off what plays the role of F, K, D and v.

- A *zero-coupon bond option*—the option at expiry T_O to buy the T_B-bond struck at K—is Black with forward equal to the forward bond price $F = P(0, T_B)/P(0, T_O)$, discount $D = P(0, T_O)$, and total volatility $v = \Sigma_P$ from (17), the accumulated log-bond volatility.
- A *caplet* is algebraically a *put* on a zero-coupon bond, so it is the put branch of (38) at that same Σ_P, scaled by the accrual factor; the worked example in §7 runs exactly this.
- A *European swaption* is a Jamshidian portfolio of such bond options (Appendix E): the coupon bond crosses the strike at a single critical rate, the swaption splits into one bond

option per cashflow, and each piece is again (38) with its own strike and the same engine (17).

So the entire linear-rates options book—bonds, caps, floors, European swaptions—reduces to evaluating (38) with the right (F, K, D, v). The model's only job is to supply the total volatility v, and in the one-factor model it supplies it in closed form.

D Anatomy of a vanilla swap and swaption

The body treats a swap purely as a difference of bonds (§7); this appendix fills in the practical mechanics behind that picture—how the payments are scheduled, when the floating rate is fixed, and how payer and receiver swaps and their options differ. None of it changes the pricing, but it is what a trader means by the words.

The two legs. A *vanilla* interest-rate swap is an agreement to exchange, on a fixed schedule and for an agreed term, two streams of interest on the same notional N in a single currency. One stream, the *fixed leg*, pays a constant rate K agreed at the outset. The other, the *floating leg*, pays a short-term reference rate—historically a LIBOR fixing, now an overnight rate such as SOFR or ESTR compounded over each period—that is reset every period. The notional itself is never exchanged; only the interest is, so a swap is a pure bet on the level of rates rather than a financing. "Vanilla" pins down all the boring choices: a single currency, a constant notional, a regular schedule, and a floating rate with no spread, cap, or exotic feature.

The schedule: accrual, fixing, payment. Fix a start date T_0 and period-end dates $T_0 < T_1 < \cdots < T_n$. The interval $[T_{i-1}, T_i]$ is the i-th *accrual period*, of length τ_i measured in a market day-count convention (e.g. Act/360 on the floating leg, 30/360 on the fixed). Three distinct dates govern each floating payment, and conflating them is the usual source of confusion (Figure 34):

- the *fixing* (or *reset*) date, at the *start* of the period T_{i-1}, when the reference rate R_i for that period is observed;
- the *accrual period* $[T_{i-1}, T_i]$, over which interest builds at that rate;
- the *payment* date, at the *end* T_i, when the cash $N\tau_i R_i$ actually changes hands.

So the floating leg fixes *in advance* and pays *in arrears*: the rate is known one period before the money moves. The fixed leg needs no fixing—its rate is K throughout—and pays $N\tau_i K$ on the same dates T_i (often on a coarser schedule, e.g. annual fixed against quarterly floating). On each payment date only the *net* of the two coupons is exchanged.

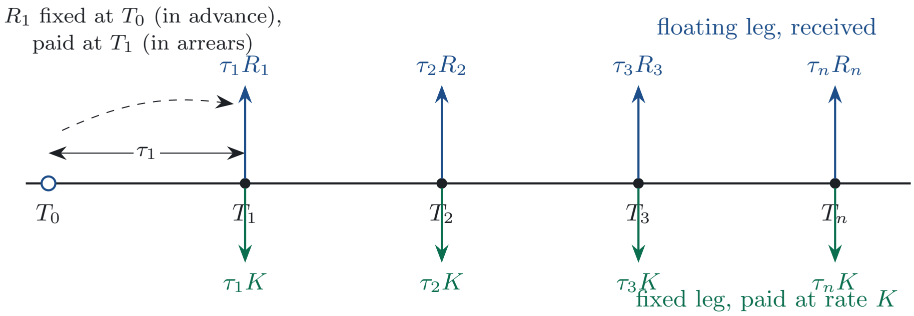

Figure 34: The schedule of a vanilla swap. Each accrual period $[T_{i-1}, T_i]$ has length τ_i. The floating rate R_i is *fixed in advance*, at the start T_{i-1} of its period, and *paid in arrears*, at the end T_i; the fixed leg pays $\tau_i K$ on the same dates. The notional is not exchanged, and on each date only the net of the two coupons changes hands. This fix-in-advance schedule is exactly what lets the floating leg telescope to $P(t, T_0) - P(t, T_n)$.

The schedule is what makes the floating leg collapse. The period-i floating coupon, with R_i the simply-compounded rate fixed at T_{i-1}, is worth $P(t, T_{i-1}) - P(t, T_i)$ in present value—the deposit-rolling argument of §7—and the sum telescopes to $P(t, T_0) - P(t, T_n)$, with no reference to the path of rates. The schedule does all the work.

Payer versus receiver. The two sides are mirror images, named for what they do with the *fixed* leg. A *payer* swap *pays* fixed and *receives* floating; a *receiver* swap *receives* fixed and *pays* floating. With $A(t) = \sum_i \tau_i P(t, T_i)$ the annuity and $S(t)$ the forward swap rate (18), their values are

$$V_{\text{pay}}(t) = N\,A(t)\,\big(S(t) - K\big), \qquad V_{\text{rec}}(t) = N\,A(t)\,\big(K - S(t)\big) = -\,V_{\text{pay}}(t),$$

equal and opposite, and both zero at the par rate $K = S(t)$. A payer gains when rates *rise*—it has locked in paying the old, lower fixed rate while receiving the new, higher floating—so it is the rate-market analogue of a long position; a receiver gains when rates *fall* (Figure 35). Because they are negatives of each other, entering a payer is precisely unwinding a receiver on the same schedule.

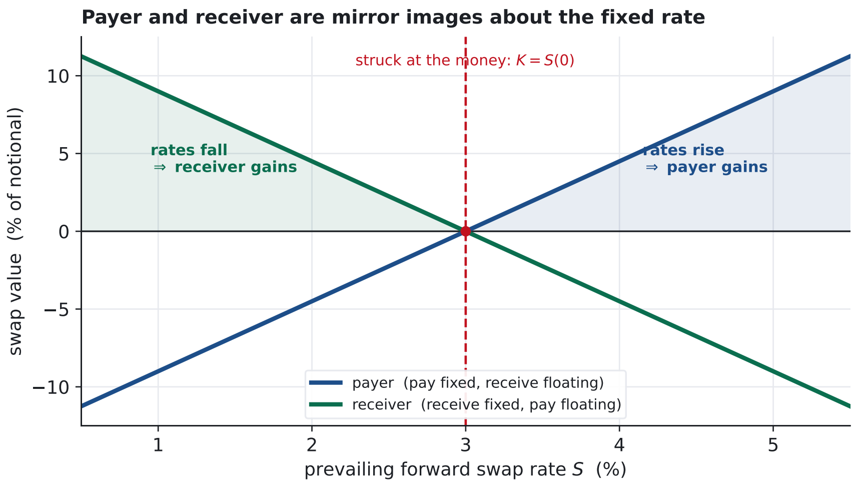

Figure 35: Payer and receiver swaps are mirror images about the fixed rate. Holding the contracted fixed rate K and the schedule, the value of each leg is the annuity times the gap between the prevailing forward swap rate S and K: the payer (blue) gains as rates rise above K, the receiver (green) as they fall below it, and both are worth nothing at the par rate $S = K$—which is where the swap is struck at the money.

From swap to swaption. A *swaption* is an option to enter a swap. A European *payer swaption* with expiry T_O and strike K is the right, not the obligation, to enter at T_O a payer swap with fixed rate K running over $T_O = T_0 < T_1 < \cdots < T_n$ (Figure 36). The option's expiry coincides with the swap's start: a "5y into 10y" payer swaption is the right, five years from now, to start paying fixed for ten years. At expiry the holder compares the par swap rate then prevailing, $S(T_O)$, with the strike and exercises only if the swap has positive value, so the payoff is

$$N\, A(T_O)\, \big(S(T_O) - K\big)^+,$$

the annuity times the in-the-money amount—which is exactly the $\big(1 - C(T_O)\big)^+$ put on the fixed-leg coupon bond of §7, valued by Jamshidian's decomposition (Appendix E). A *receiver swaption* is the matching right to receive fixed, paying $N\, A(T_O)\, \big(K - S(T_O)\big)^+$; the two obey put–call parity against the forward swap—payer minus receiver is the forward swap itself, the swap-market echo of the cap−floor=swap identity of §7. Settlement is either *physical* (the parties actually enter the swap) or *cash* (its mark-to-market value is paid); in a one-factor Gaussian model the distinction does not affect the price.

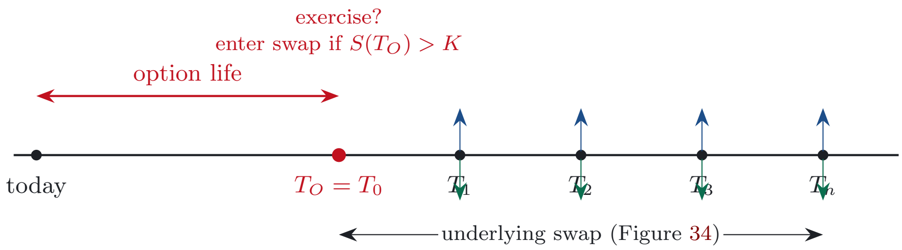

Figure 36: A European swaption in time. The option lives from today to its expiry T_O, which is also the start date T_0 of the underlying swap. At T_O the holder exercises only if the swap is in the money—for a payer, if the prevailing swap rate $S(T_O)$ exceeds the strike K—and then receives the swap, whose floating (blue) and fixed (green) exchanges run on to T_n. The expiry/tenor pair is what the market quotes as, e.g., "5y into 10y".

E Jamshidian's decomposition for swaptions

A European swaption is an option on a coupon bond—the fixed leg of the underlying swap, viewed as a stream of payments c_i at dates T_i, plus notional. Exercise is worthwhile when the coupon bond's value $\sum_i c_i\, P(T_O, T_i)$ crosses the strike at the option expiry T_O. In a *one-factor* model every $P(T_O, T_i)$ is a monotone (decreasing) function of the single state r_{T_O}, so the coupon bond is monotone in r_{T_O} too, and there is a *unique* critical rate r^* at which it equals the strike. Jamshidian's observation [5] is that one can then set, for each cashflow date, a strike $K_i = P(T_O, T_i)\big|_{r_{T_O} = r^*}$—the value that bond takes at the critical rate. Because all the bonds cross their respective K_i at the *same* rate r^*, the payoff of the option on the *portfolio* equals the sum of the payoffs of options on the *individual* bonds struck at the K_i:

$$\Big(\sum_i c_i\, P(T_O, T_i) - K\Big)^+ \;=\; \sum_i c_i\, \big(P(T_O, T_i) - K_i\big)^+. \tag{39}$$

Each term on the right is a European option on a zero-coupon bond, priced in closed form by (17). The swaption is thus an exact portfolio of bond options. The argument needs the common crossing point r^*, which exists only because one state variable orders all the bonds together—precisely the monotonicity that two factors destroy, which is why Part II's swaptions fall back on a one-dimensional numerical integration instead.

F Solving the OU process: mean and variance of r_t

This appendix derives (11) in full from the decomposition (10). The short rate is $r_t = x_t + \alpha(t)$ with $\alpha(t)$ deterministic, so $\mathbb{E}[r_t] = \mathbb{E}[x_t] + \alpha(t)$ and $\mathrm{Var}[r_t] = \mathrm{Var}[x_t]$: the deterministic shift moves the mean but contributes nothing to the spread. Everything therefore reduces to the mean and variance of the zero-mean Ornstein–Uhlenbeck (OU) factor

$$\mathrm{d}x_t = -a\,x_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \qquad x_0 = 0, \qquad a > 0. \tag{40}$$

We first solve (40) exactly, then read off the moments two independent ways each (a closed-form route and a differential-equation route, as a cross-check), and finish with the full distribution, the autocovariance, and the exact transition law used to simulate the process.

F.1 The integrating factor, and why it has no Itô correction

Equation (40) is *linear* in x_t: the drift $-a\,x_t$ is proportional to the state and the diffusion coefficient σ is constant. The standard device for a linear ODE—multiply by an integrating factor that turns the left-hand side into an exact derivative—works here too, and the right factor is e^{at}, chosen so that its growth rate a cancels the mean reversion.

The one subtlety is that we are differentiating a *product* of a deterministic function e^{at} and a stochastic process x_t, so in principle we must use Itô's product rule,

$$\mathrm{d}\big(e^{at}\,x_t\big) = e^{at}\,\mathrm{d}x_t \;+\; x_t\,\mathrm{d}\big(e^{at}\big) \;+\; \mathrm{d}\big\langle e^{a\cdot}, x\big\rangle_t. \tag{41}$$

The last term is the quadratic covariation between e^{at} and x_t. But e^{at} is a smooth, finite-variation function—it has no Brownian part—so its quadratic covariation with anything is zero, $\mathrm{d}\langle e^{a\cdot}, x\rangle_t = 0$, and (41) collapses to the *ordinary* product rule. With $\mathrm{d}(e^{at}) = a\,e^{at}\,\mathrm{d}t$,

$$\mathrm{d}\big(e^{at}\,x_t\big) = e^{at}\,\mathrm{d}x_t + a\,e^{at}\,x_t\,\mathrm{d}t = e^{at}\big(\underbrace{-a\,x_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t}_{\mathrm{d}x_t}\big) + a\,e^{at}\,x_t\,\mathrm{d}t = \sigma\,e^{at}\,\mathrm{d}W_t, \tag{42}$$

the two $a\,e^{at}\,x_t\,\mathrm{d}t$ terms cancelling exactly—which is the whole point of the factor e^{at}.

F.2 The explicit solution

Integrate (42) from 0 to t. The left-hand side telescopes to $e^{at}x_t - e^0x_0 = e^{at}x_t$ since $x_0 = 0$:

$$e^{at}\,x_t = \sigma\int_0^t e^{as}\,\mathrm{d}W_s \qquad \Longrightarrow \qquad x_t = \sigma\int_0^t e^{-a(t-s)}\,\mathrm{d}W_s. \tag{43}$$

This is the exact solution. Read it as a weighted sum of all the shocks the process has received: the increment $\mathrm{d}W_s$ that arrived at time s is still present at time t, but discounted by $e^{-a(t-s)}$—a shock $\tau = t-s$ years old has faded to a fraction $e^{-a\tau}$ of its original size. Recent shocks dominate; old ones have decayed away. This exponential forgetting *is* mean reversion, written out explicitly.

Consistency check. It is worth confirming (43) really solves (40). Write $x_t = e^{-at}M_t$ with the martingale $M_t = \sigma\int_0^t e^{as}\,\mathrm{d}W_s$, so $\mathrm{d}M_t = \sigma e^{at}\,\mathrm{d}W_t$. Applying the product rule once more (again e^{-at} is finite-variation, so no correction term),

$$\begin{aligned}\mathrm{d}x_t &= -a\,e^{-at}M_t\,\mathrm{d}t + e^{-at}\,\mathrm{d}M_t\\ &= -a\,x_t\,\mathrm{d}t + e^{-at}\,\sigma e^{at}\,\mathrm{d}W_t\\ &= -a\,x_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t,\end{aligned} \tag{44}$$

which is (40), and $x_0 = e^0M_0 = 0$. The solution is verified.

F.3 The mean

Route 1 — the martingale property. The Itô integral $\int_0^t f(s)\,\mathrm{d}W_s$ of a deterministic (more generally square-integrable adapted) integrand is a martingale started at 0, hence has mean 0: each Brownian increment $\mathrm{d}W_s$ has mean zero and is independent of the past, so it adds nothing in expectation. Taking expectations of (43),

$$\begin{aligned}\mathbb{E}[x_t] &= \sigma \int_0^t e^{-a(t-s)}\,\mathbb{E}[\mathrm{d}W_s] \\ &= 0.\end{aligned} \tag{45}$$

Route 2 — the moment ODE. Equivalently, take expectations of the SDE (40) directly. With $m(t) := \mathbb{E}[x_t]$ and $\mathbb{E}[\mathrm{d}W_t] = 0$,

$$\mathrm{d}\,\mathbb{E}[x_t] = -a\,\mathbb{E}[x_t]\,\mathrm{d}t \qquad \Longrightarrow \qquad m'(t) = -a\,m(t), \quad m(0) = 0, \tag{46}$$

a linear ODE whose only solution is $m(t) \equiv 0$. Either way, since the shift is deterministic,

$$\boxed{\begin{aligned}\mathbb{E}[r_t] &= \mathbb{E}[x_t] + \alpha(t) \\ &= \alpha(t).\end{aligned}} \tag{47}$$

F.4 The variance, by Itô's isometry

Because x_t has zero mean, its variance is just $\mathbb{E}[x_t^2]$. The tool is Itô's *isometry*: for a deterministic square-integrable f,

$$\mathbb{E}\left[\Big(\int_0^t f(s)\,\mathrm{d}W_s\Big)^2\right] = \int_0^t f(s)^2\,\mathrm{d}s. \tag{48}$$

The reason is the independence of disjoint Brownian increments. Expanding the square produces a double integral over (s,u) of $f(s)f(u)\,\mathrm{d}W_s\,\mathrm{d}W_u$; in expectation the off-diagonal terms vanish, because for $s \neq u$ the increments are independent with mean zero, $\mathbb{E}[\mathrm{d}W_s\,\mathrm{d}W_u] = 0$, while on the diagonal $\mathbb{E}[(\mathrm{d}W_s)^2] = \mathrm{d}s$. Only the diagonal survives, leaving $\int_0^t f(s)^2\,\mathrm{d}s$. Applying (48) to (43) with $f(s) = \sigma\,e^{-a(t-s)}$,

$$\begin{aligned}\mathrm{Var}[x_t] &= \sigma^2 \int_0^t e^{-2a(t-s)}\,\mathrm{d}s \\ &= \sigma^2 e^{-2at} \int_0^t e^{2as}\,\mathrm{d}s \\ &= \sigma^2 e^{-2at}\,\frac{e^{2at}-1}{2a} \\ &= \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big),\end{aligned} \tag{49}$$

where the middle step pulls the constant e^{-2at} out of the integral and the remaining $\int_0^t e^{2as}\,\mathrm{d}s = (e^{2at}-1)/2a$ is an ordinary integral.

F.5 The variance again, by a moment ODE

As an independent check, derive the same variance without the explicit solution. Apply Itô's formula to x_t^2, where now the second-order term *does* contribute because $x \mapsto x^2$ is nonlinear and x_t carries a Brownian part of volatility σ:

$$\begin{aligned}\mathrm{d}\big(x_t^2\big) &= 2x_t\,\mathrm{d}x_t + \mathrm{d}\langle x\rangle_t \\ &= 2x_t\big(-a\,x_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t\big) + \sigma^2\,\mathrm{d}t \\ &= \big(\sigma^2 - 2a\,x_t^2\big)\,\mathrm{d}t + 2\sigma x_t\,\mathrm{d}W_t.\end{aligned} \tag{50}$$

Take expectations; the $\mathrm{d}W_t$ term is a mean-zero martingale increment and drops out. Writing $v(t) := \mathbb{E}[x_t^2] = \mathrm{Var}[x_t]$,

$$v'(t) = \sigma^2 - 2a\,v(t), \qquad v(0) = 0. \tag{51}$$

This is again first-order linear: multiply by the integrating factor e^{2at} to get $\big(e^{2at}v\big)' = \sigma^2 e^{2at}$, integrate from 0 to t, and divide back:

$$e^{2at}v(t) = \frac{\sigma^2}{2a}\big(e^{2at} - 1\big) \qquad \Longrightarrow \qquad v(t) = \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big), \tag{52}$$

matching (49) exactly. The ODE form (51) also makes the behaviour transparent: variance is *injected* at the constant rate σ^2 by the random shocks and *bled off* at rate $2a\,v$ by the mean reversion, and the two balance at the stationary level $v = \sigma^2/2a$. Since the shift is deterministic,

$$\boxed{\begin{aligned}\mathrm{Var}[r_t] &= \mathrm{Var}[x_t] \\ &= \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big).\end{aligned}} \tag{53}$$

F.6 The full distribution

More than the first two moments is available almost for free. The solution (43) is a *Wiener integral*—a stochastic integral of a deterministic function—and such an integral is a limit of sums of independent Gaussian increments, hence itself *Gaussian*. So x_t is normally distributed, and therefore so is the short rate:

$$r_t \sim \mathcal{N}\bigg(\alpha(t),\ \frac{\sigma^2}{2a}\big(1 - e^{-2at}\big)\bigg). \tag{54}$$

This is the fact underlying everything analytic in Part I: it is precisely because r_t (and every linear functional of it, including the log bond price) is Gaussian that caps, floors and swaptions admit the closed-form Black-style prices of Section 7—and also why the rate can in principle go negative, the normal distribution having support on the whole line.

F.7 The autocovariance

The same machinery gives the covariance between the factor at two different times $s \le t$, which is what governs how rates at different horizons co-move. Split the later integral (43) at s,

$$x_t = \underbrace{e^{-a(t-s)}\,x_s}_{\text{decayed past}} + \underbrace{\sigma\int_s^t e^{-a(t-u)}\,\mathrm{d}W_u}_{\text{new shocks, independent of } x_s}, \tag{55}$$

and note the second piece is built from increments after s, hence independent of x_s and of mean zero. Therefore

$$\begin{aligned}\mathrm{Cov}(x_s, x_t) &= e^{-a(t-s)}\,\mathbb{E}[x_s^2] \\ &= e^{-a(t-s)}\,\frac{\sigma^2}{2a}\big(1 - e^{-2as}\big) \\ &= \frac{\sigma^2}{2a}\Big(e^{-a(t-s)} - e^{-a(t+s)}\Big).\end{aligned} \tag{56}$$

For large s and t with the gap $\tau = t - s$ fixed (the stationary regime) this tends to $\frac{\sigma^2}{2a}\,e^{-a\tau}$: correlations between the factor at dates τ apart decay exponentially with a timescale $1/a$. This is the single-factor root of the perfect-correlation defect of Section 11—one OU factor imposes one rigid exponential correlation structure on the whole curve.

F.8 The transition law used in simulation

The split above, kept general, is also the exact recipe for stepping the process forward in a simulation. Conditional on x_s, the factor at a later time t is Gaussian with

$$\begin{aligned} \mathbb{E}[x_t \mid x_s] &= e^{-a(t-s)} x_s, \\ \operatorname{Var}[x_t \mid x_s] &= \frac{\sigma^2}{2a}\big(1 - e^{-2a(t-s)}\big). \end{aligned} \tag{57}$$

Writing the step size $\Delta = t - s$, this means one may advance the process over a step of *any* size with no discretisation error at all,

$$x_{s+\Delta} = e^{-a\Delta}\, x_s + \sigma\sqrt{\frac{1 - e^{-2a\Delta}}{2a}}\, Z, \qquad Z \sim \mathcal{N}(0,1) \text{ i.i.d.}, \tag{58}$$

because (58) reproduces the exact conditional moments (57) (and the process is Gaussian, so the first two moments determine the law completely). This is the scheme used to generate the short-rate paths in Figures 8–9: unlike an Euler step, which would only be approximate, (58) is exact for any Δ, a direct dividend of the model's linear-Gaussian structure.

F.9 Limiting behaviour

As $t \to \infty$ the factor $(1 - e^{-2at}) \to 1$, so the variance saturates at the stationary value $\sigma^2/(2a)$ and the distribution of x_t settles to its invariant law $\mathcal{N}(0, \sigma^2/2a)$—the fan in Figure 8 reaching a stable width. Starting from $x_0 = 0$, the variance covers half its long-run value at the *half-life* $t_{1/2} = \ln 2\,/\,(2a)$. A fast mean reversion (large a) reaches this steady width sooner *and* makes it narrower—both because $\sigma^2/2a$ shrinks with a—which is the precise sense in which a controls how wide the long-run distribution of rates is, as claimed beneath (11).

References

[1] O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

[2] J. C. Cox, J. E. Ingersoll, S. A. Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.

[3] J. Hull, A. White. Pricing interest-rate-derivative securities. *Review of Financial Studies*, 3(4):573–592, 1990.

[4] J. Hull, A. White. Numerical procedures for implementing term structure models I: Single-factor models. *Journal of Derivatives*, 2(1):7–16, 1994.

[5] F. Jamshidian. An exact bond option formula. *Journal of Finance*, 44(1):205–209, 1989.

[6] D. Brigo, F. Mercurio. *Interest Rate Models—Theory and Practice*. Springer, 2nd edition, 2006.

[7] L. B. G. Andersen, V. V. Piterbarg. *Interest Rate Modeling*. Atlantic Financial Press, 2010.

[8] A. Brace, D. Gatarek, M. Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7(2):127–155, 1997.

[9] D. Filipović. *Term-Structure Models: A Graduate Course*. Springer, 2009.